

David Griesinger
Lexicon
100 Beaver St.
Waltham, MA 02154

March 21, 1989

Practical Processors and Programs for Digital Reverberation

Abstract

Digital Reverberation is everywhere. In only eleven years it has gone from one of the first really cost-effective uses of digital technology to being indispensable to modern recording. This paper presents some of the origins of digital reverberation in the extensive earlier work at the British Broadcasting Corporation, Bell Laboratories, and elsewhere. It will discuss the advantages and disadvantages of some of the algorithms, and the hardware requirements of digital processors optimized for reverberation.

IT CAN'T BE DONE

Before discussing electronic reverberation it is worth pointing out why it is really not possible to electronically emulate natural reverberation. In practice there are two reasons, one of which, the complexity of natural impulse responses, merely makes the problem impracticably difficult. The other problem, the multitude of sources, makes it impossible. Once we have done our duty by pointing these problems out, we can continue with a good conscience to describe what we CAN do.

IMPULSE RESPONSE OF NATURAL SPACES

If we shoot a pistol at one point in a room, and record the sound pressure with microphone at another point we can capture an echogram, or the impulse response of the room. If we have a computer powerful enough to CONVOLVE recorded music with this measured impulse response we can duplicate the sound of the room exactly - just as if we had a perfect loudspeaker at the position of the pistol.

The process of convolution sounds more complicated than it is. First we digitize the impulse response and store it in the computer. To perform the convolution we take the first sample of the impulse response and multiply by the first sample of the music, and then add the product to the product of the second sample of the impulse response multiplied by the second sample of the music. We keep doing this until the impulse response has decayed enough that we needn't continue. The resulting sum is the first output sample. To find the second output sample we do the same thing, but this time multiplying the first sample of the impulse response by the SECOND sample of the music, and summing everything again shifted by one sample. Each sample of the output comes from a sum of a great many multiplications. If our sample rate is 44kHz, and the room has a 1 second reverb time, each output sample is the result of 44 thousand multiply and adds. To do these in real time

requires a machine with 2000 mega instructions per second (MIPS) capacity. For a 2 second reverb time, we need twice this, or 4000 MIPS. For stereo we need 8000 MIPS.

We could use FFT techniques to do the convolutions - but this jumps ahead a bit. For the moment lets think about what happens if we simplify the impulse response so it contains only a few multiplies.

Natural spaces have response functions of great complexity. If you record a pistol shot and play it back on an oscilloscope you get patterns similar to Figures 1 and 2. It is tempting to claim such an echograms can be characterized by just a few simple reflections. If the impulse response contains only 100 reflections we need only 100 multiply and adds -5 MIPS -a much more practical figure. Nearly all work in electronic reverberation demands such simplification. Does it work? Not really.

A few prominent reflections are easy to identify in figure 1, but if we look more closely we find each reflection is not a simple spike but a more complicated shape. The spread in the individual reflections is due to the surface roughness and frequency dispersion of natural reflectors, a roughness which can easily extend over a time period of more than a millisecond. In good rooms the walls are bumpy -and who wants to record in bad rooms?

In trying to reduce the problem to manageable size we have substituted a single 20 microsecond wide delay for the broader shape of a natural reflection. The difference in sound is large. Ironically the higher the bandwidth of the electronic reverberator the greater is the discrepancy, and the more unnatural the simple delay will sound. In addition there is a great deal of what looks like noise throughout the diagram, and at later times -see figure 2 -this noise begins to dominate the whole response. These wiggles are reflections from objects and surfaces of small size. These small reflections are vital to the perceived smoothness and diffusion in the reverberation.

MULTIPLE SOURCES

The impulse responses of rooms are chaotic, in that the impulse response depends strongly on the position of the source and the listener. Anytime you move either the microphone or the position of the sound source the impulse response changes dramatically, as does the timbre of any music convolved with this impulse response. Each musician will have their timbre modified by the room in a different way, and in addition each of them will make small movements to their instruments as they play. It is the superposition of all these changing impulse responses which we recognize as the sound of a good hall. To emulate it electronically we need a separate reverberation device for each instrument, each using an impulse response which changes slightly with time.

It is this requirement for a great multitude of input channels which limits the performance of current reverberation devices most strongly. When you have only one or two input channels the characteristic timbre of the reverberation device will always be audible, since the same timbre is applied to everything. We can do nothing about this problem, so lets ignore it for the moment. We should not complain, since the coloration which must result

from a lack of multiple sources can mask colorations which arise from trying to generate reverberation with a practical processor.

LET'S DO IT ANYWAY -WITH RECIRCULATION

The process of convolution described above is a method of generating reverb which uses FIR or finite impulse response filters. FIR's are all limited by the number of delays you can sum in a practical processor. In digital filters involving long delay times or low frequencies it is usually better to use feedback or recirculation to get the response you want. Recirculated delays form IIR or infinite impulse response filters. Can we build a reverberator with IIR filters? The idea is promising, since such filters have inherently an exponentially decaying impulse response, and one multiply and one delay can account for many reflections.

REVERBERATION WITH ACOUSTIC AND MAGNETIC DELAYS

This is not a new problem. Recirculation of delayed sound has been used to emulate reverberation for many years. I will pick up the story with the work of Axon, Gilford, and Shorter at the BBC in the 1950's. They used acoustic delays (sound traveling in a pipe) and magnetic recorders to extensively study the reverberant behavior of recirculation.

The simplest recirculator is a single delay with feedback. Figure 3 -from Schroeder (1962). If the attenuation (the coefficient of the multiply) is g , the attenuation in decibels each time through the loop is $-20 \cdot \log(g)$. Thus if the delay line has length t , the reverb time (the time it takes for a sound to decay 60dB) is

$$T = t * 60 / 20 \log(1/g) = 3t / \log(1/g) \text{ seconds.}$$

The impulse response and the tone response are given in the figure. The tone response is the familiar "comb filter", where the sharpness of the teeth depends on the attenuation g . If g is nearly one -to get a long reverb time - the teeth are very sharp indeed. Input frequencies which fall on the peaks of the response will reverberate. All others decay away quickly. Any variation in the frequency response of the delay line will cause a large change in the reverb time of the peaks as a function of frequency.

Axon et al point out that as a reverberator this system has two shortcomings. 1. The impulse response is not dense enough -individual echos can be heard as flutter, and 2. the timbre is very strongly colored by the fundamental resonance of the filter and its harmonics -which are the only frequencies which reverberate! If we make the delay longer the fundamental frequency becomes lower so the modal density (the number of comb teeth per octave) becomes higher, but the time density of reflections becomes lower.

We can add some more taps to the delay, while keeping the feedback confined only to the last Tap. Figure 4. This gives the impulse response shown. This arrangement is promising. We can increase the density of the impulse response simply by adding taps. Notice however that the impulse response has constant density with time - once the pattern of taps is

established it remains constant until the sound has decayed entirely.

How about using two taps to derive the feedback? Figure 5. Note that now the impulse response becomes more dense as time progresses - much more similar to the natural case. Unfortunately the amplitudes of the individual reflections fluctuate wildly. The same or worse fluctuation occurs in the tone response. In fact as the feedback is increased a few prominent frequencies will have very long reverb times, or even oscillate. Such a reverberator sounds very similar to a PA system running just below feedback. Axon and Gilford concluded that feedback should always come from a single tap, because the coloration and the stability of such a system is always predictable. Later work at the BBC (see Spring and Gilford) investigated the results of randomly switching or blending between two separate taps for feedback, with more natural results.

Axon et al also tried putting several combs such as figure 4 in parallel, and this was the basis of a magnetic reverberation system built at the BBC. Several parallel comb filters of different lengths with multiple taps involving non commensurate lengths have many advantages over simpler systems. The density of modes in the tone frequency response can be made large, as can the density of reflections in the impulse response. Again later experiments at the BBC with randomly varying the position of the delay taps resulted in significant improvements. Note that for any parallel system the number of reflections in any interval of time is proportional to the number of sections. Achieving twice the impulse density requires twice as much computation. We really would like to put reverberators in series.

DIGITAL REVERBERATION

Although the analysis of Axon and Gilford was insightful their hardware was limited and expensive. Electronic reverberation requires delay units with excellent signal to noise ratio, low distortion, and completely flat frequency response. These are properties of digital systems. It is no accident that the first digital audio product was a delay line, and the first real-time use of digital signal processing in audio was for reverberation. (see Blesser and Bader).

M.R. Schroeder was the first to try making artificial reverberation through computer manipulation of digitized sounds. Schroeder's computer at Bell Labs was capable of working only in non-real time on relatively short samples of sound. However he made an exceedingly important contribution by recognizing the usefulness of IIR filters with all-pass or flat tone frequency response in the synthesis of reverberation. Schroeder pointed out that if you add a negative feed forward path around a comb filter the frequency response can be made flat. Figure 6.

Note that although the steady state tone response is flat, the phase response, and the response to a rapidly varying signal such as music, is not flat at all. However all pass filters have a much smaller effect on timbre than a comb of similar delay and reverb time. The major advantage of all-pass filters is that they may be wired in SERIES. Ordinary comb filters placed in series sound very bad, since only tones which just happen

to resonate in all of them will pass through. All pass filters do not have this problem -as many as you want can be placed in series and they will still pass all frequencies. The impulse response can become quite complex, since the number of impulses in any given time is multiplied in each section.

Schroeder found that if sound were passed through two or more all pass filters in series after going through a network of comb filters in parallel a much denser impulse response results. Figure 7.

LATER WORK

Schroeder's algorithm sounds amazingly good, especially considering that he did not have a chance to listen to it with extended pieces of music. Later workers extended his ideas in several directions, although most of this work, including our own, remains unpublished. In the published literature Moorer and Gerzon have shown different types of all pass filters. Moorer suggested adding filtering to the combs to simulate air absorption (and add badly needed smoothing to the sharp reflections). He also suggests adding a FIR filter to the start of the reverberator in order to simulate the relatively discrete reflections at the start of reverberation. The output of the FIR filter feeds the 'reverberation output directly, and some or all of the taps are used as inputs to a system of parallel combs which generates the more complex and diffuse part of the decay. The care with which the taps are chosen has a large effect on the sound quality of the result.

Stautner's reverberator is an example of a design with multiple feedback taps. As Axon et al found out this configuration is not always stable, but it does provide a more natural build up of impulse density. Stautner was able to control some of the unpleasant resonant behavior by using time varying delays and low values of feedback.

There is also a method developed by J.O. Smith at Stanford which models reverberation through simulating reflections inside multiply connected wave guides. This technique allows a build-up in echo density while maintaining stability.

In addition at Lexicon we tailor the shape and the spread of the beginning of the diffused reverberation. As will be pointed out below, the shape of the first 15dB of decay is vital to the sound quality.

Although the basic elements of reverberant filters are well known, we want to emphasize that making an acceptable reverberator is still very difficult. Wood and strings for a violin are easy to buy, but making an excellent instrument demands skill, experience, and patience. Due to the ease with which mathematical algorithms may be copied most of the people who work in this field are compelled to keep the details of their knowledge to themselves.

Several tests have been developed to help determine reverberation quality.

1. Plots of reflection density versus time. My preferred way is to count the number of echos within 20 dB of the largest echo, using a sliding window size of about 20ms. This

function can be very different for different algorithms.

2. Plots of energy (integrated impulse response) as a function of time.
3. Plots of the autocorrelation function of the reverb output when excited by a click. Before doing the autocorrelation the reverberator output should be multiplied by an exponentially increasing function, so the total level remains relatively constant as the sound decays. This plot is very good at showing repetitive behavior in the reverberant decay. See K.A. Law.
4. Sonograms - plotting the frequency content as a function of time in a reverberator as the sound decays. Again the reverberator output should be multiplied by an increasing exponential. These plots are very good at detecting frequencies which have longer reverberation times than others. This technique was developed by MR. Schroeder.
5. Plots of the reverberation time of each mode as a function of frequency. In general different modes of a reverberator can decay at different rates. At the end of the decay only the longest lived modes remain, and the sound can be thin and metallic.
6. Plots of the steady-state amplitude response as a function of frequency. This is helpful in determining the modal density, and can point out if there are regular gaps in the reverberant modes.
7. All these tests can be tedious. The best test is to simply listen to the algorithm, using a combination of clicks, boings, and music.

HOW DOES IT SOUND?

Lets start with "Acoustic" sounds, such as room sound and concert halls. It is important to realize that the best reverberation for recorded sound may be different from the best live concert hall acoustics. Reverberation in halls and rooms has considerable spatial and directional information which is poorly reproduced by conventional stereo recordings, and this has a major impact on understanding the relation between architectural acoustics and recorded acoustics.

ACOUSTIC SPACES

A few general suggestions can be made from the impulse response of actual halls: First, the overall complexity or density of the echogram is important, especially when comparing different types of artificial reverberation. Schroeder proposed that the echo density should be at least 1000 reflection/second if the reverberation was to sound natural. A quick look at figure 1 shows this hall easily meets this criterion, almost from the very start. In addition the echos in figure one are probably not a precise copy of the original impulse. Since many or most reflecting surfaces are small and/or rough there will be considerable smearing of the reflections in time, corresponding to various kinds of filtering in frequency. Digital reverberation has difficulty providing enough density, and even more

difficulty smearing each reflection differently in frequency.

Second, the general shape of the energy content in the first 200ms is important. Figure 2 shows the impulse response of a hall at different source-receiver distances. Note that reverberant energy comes in several distinct groups, but in general builds to a maximum 30 to 40 ms after the direct sound. The patterns formed by these fluctuations are very important to the perceived size and spaciousness of the room. In fact, simply changing the time scale of the echos in the echogram can vary the apparent size from a large concert hall to a small room. The apparent size and reverb time might be the same in different seats, but the clarity of the music might be quite different. Artificial reverbs with these echograms might be perceived as having different spatial properties while the music was playing, but might be quite similar when the music stopped.

Third, the effective reverb time while the music is playing (which is most of the time) is determined by the initial rate of decay of the impulse response, say from max to about -15dB. The classical reverberation time as formulated by Sabine is usually of secondary importance.

Fourth, many of the desirable properties of halls, described by such words as spaciousness, spatial impression (SI), and envelopment, depend on directional properties of the reflections in the first 200ms, and it is not clear how these properties translate to artificial reverberation from stereo loudspeakers. However, it seems to be very important that a stereo reverberator should have uncorellated impulse responses at the two outputs.

Fifth, reverberation for recording differs from reverberation for live listening in the importance of early (10 to 50 ms) reflections. In halls early reflections allow the musicians on stage to hear each other, and add to the loudness and intimacy of the sound to the audience. However, most engineers find that these reflections, usually from the ceiling, floor and side walls of the stage enclosure, can seriously muddy recorded music. Engineers try to place sound sources and microphones to minimize these reflections, and the common practice of moving the orchestra off the stage onto the floor of a concert hall for a recording is a good example.

Sixth, the reverberant decay at time delays greater than 200ms is important too, there are also complications here. Decay in actual rooms may not be either exponential or smooth. Many rooms exhibit a longer reverb time at the end of the decay than at the beginning, or may have an obvious pulsation as sound bounces from the back to the front which never really smoothes out. Depending on the algorithm used, an artificial reverberation device may have other problems as the sound decays, either insufficient diffusion, unevenness, or a metallic quality. Some electronic reverberation devices allow the user to select unnatural final decay rates, either much shorter than you would expect from the apparent decay rate in the first 15 dB of decay (the running reverb), or much longer. These adjustments (especially making the stopped decay short) are very useful and popular.

Seventh, the time density of reflections is very important to the subjective quality of the reverb, especially if sharp percussive sounds are present in the music. As discussed before,

in natural acoustics a reflection is almost never really completely discrete. In good spaces walls are seldom perfectly flat or parallel, and reflections are always somewhat smeared out in time. Unfortunately, in a digital reverberation device reflections are usually very discrete. If the device is excited by a sharp click, the resulting multiple clicks can sound very unnatural, giving a grainy, almost distorted sound to drums or woodblocks. The density of echos necessary to remove this graininess depends both on the bandwidth of the machine and the abilities of the listener, but can be much higher than Schroeders proposal, up to 10,000 echos/second or beyond. The higher the bandwidth of the reverberation device the greater the density must be. Techniques used to add diffusion (density) to the electronic reflections add problems of their own, such that a device set for high diffusion might sound good on one type of music, and one set for low diffusion might sound better on another type of music.

In natural spaces the time density of reflections starts at some value determined by the complexity of the wall surfaces, and builds up at a rate which depends on the size and shape of the space. Even the initial reflections are smeared enough to lose the sharp grainy quality of a precise digital delay. By the time the reverberation is 15 dB down the density will most likely be so high individual reflections will be impossible to distinguish by ear . Most digital reverberation algorithms used today do not increase the time density fast enough, nor do they continue building the density as the sound decays.

CHAMBERS AND PLATES

To emulate the sound of a good reverberation chamber or a reverberation plate the algorithm should concentrate on a very rapid build-up of impulse density, followed by a smooth exponential decay. High density can be achieved at the cost of some metallic coloration, which may be desirable in this application.

HARDWARE REQUIREMENTS

High quality reverberation demands high quality digital hardware. Fortunately digital converters of good quality are now easily available, as are the techniques needed for evaluating and testing them. The requirements on the processor used for the reverberation algorithm are less obvious.

Processing power:

The number of arithmetic operations needed for reverberation depend on the bandwidth of the processor. The lower the bandwidth the fewer delay taps need be processed to yield adequate density in the impulse response. Our early work showed that about 100 total delay taps would be adequate with a bandwidth of 8kHz. To get the same results at higher bandwidth more taps must be used, and/or a more complex algorithm. Thus we need at least 100 multiplies and adds in each sample period. Reverberation is interesting in that it does not require great precision in the coefficients used for each delay tap. For some parts of the algorithm as few as 3 bits (plus sign) are sufficient.

Data memory: I.

The larger the data memory the easier it becomes to generate long reverb times. If a reverb time of 2 seconds is being created using a total delay length of only 100ms the recirculation coefficients must be quite close to one. Obvious repetitive behavior and coloration will result. It is better to have more data memory. Up to 1/2 second is needed for standard concert halls and rooms, and up to 2 seconds of memory can be used to good advantage.

Note that the modal density is directly proportional to the total length of memory used in the algorithm. If an algorithm uses only 100ms of memory the frequencies where there is any reverberation must be separated by at least 10Hz.

This requirement for large data memory means that a high quality processor must contain from between 32k and 64k words of data memory. If dynamic memory is used its cycle time will limit the number of taps you can create in one processing block. Static memory is faster, but expensive.

Memory Access:

Digital filters used for reverberation synthesis tend to be memory access intensive. If we were simply adding up delay taps we would need an access to the main data memory for each tap we sum. Many digital signal processing (DSP) chips on the market today access external memory inefficiently, and thus are not good candidates for reverberation processors.

Memory Management:

The basic element of digital reverberation is the delay line. It can be useful to add memory management to external memory stores so that it appears as a continuous shift register to the processor. This means that the processor need not subtract a unit from each data memory address after each sample period.

Number of Bits:

Since reverberation is computation intensive there is significant build-up of noise each time a multiplication or an addition is made, unless the operations are carried out with greater precision than the usual 16 bits. 20 bits of accuracy in intermediate calculations is usually enough, unless very long reverb times are being calculated. It is also useful to have an input converter of very high quality, since noise in the input will build up in the reverberator as sound is recirculated.

Chips:

The first reverberators on the market used custom processors built from discrete parts in order to achieve enough power. 10 years ago 5 MIPS was a big deal. Nowadays there are several DSP processors which are capable of credible reverberation. We find at Lexicon that

custom processors which are designed specifically for the task still give the best ratio of price to performance.

Convolution through Fourier Analysis:

It is possible to build a machine which can do the convolution described at the beginning of this paper through Fast Fourier Techniques (FFTs.) Such a machine would be expensive. It also would have a major problem. To do the convolution a block of digital data at least the length of the reverb time must be transformed, summed with the transformed impulse response, and transformed back. The two transforms can not be done until the whole block has been collected. Thus the minimum delay of sound in passing through the reverberator would be twice the reverb time, and a practical machine would need at least three times. Such a delay could be accommodated in a fully digital console, but it would be very inconvenient. The results might be worth the wait. Moorer has evaluated reverberation generated with convolution, and reports that it sounds very good.

CONCLUSIONS

Although the techniques and algorithms described in this paper are relatively simple, creating useable reverberation with them is still largely a matter of clever thinking, careful listening, and a great deal of patience. As time goes on digital reverberators have steadily improved in sound quality and gone down in price. The performance of the best of them may be near the limits imposed by having only two outputs and two inputs. I feel the future of digital reverberation lies in removing these limitations. Already devices have appeared for the home which realistically reproduce acoustic environments through multiple loudspeakers. Multiple input devices may yet appear.

REFERENCES:

Axon, P.E., Gilford C.L.S., Shorter D.E.L.; Artificial Reverberation Proc. Instn elect. Engrs, 102B 5, pp 624-42 (1955)

Schroeder M.R., Logan B.F.; 'Colorless' Artificial Reverberation JAES v 9 #3 pp 192-197 (1961)

Schroeder M.R.; Natural Sounding Artificial Reverberation JAES 10(3): 219-223 (1962)

Gerzon M.A. Synthetic Stereo Reverberation Studio Sound 13 pp 632-635 (1971) and 14 pp 209-214 (1972)

Spring N.F., Gilford C.L.S.; Artificial Reverberation BBC Engineering 97 Mar. 1974 pp 28-32 -describes work with random delays

Blessner B., Bader K.O. Zaorski R.; A Real Time Digital Computer for Simulating Audio Systems JAES V 23 #9, P. 698 (1975)

Gerzon M.A. Unitary (Energy Preserving) Multichannel Networks with Feedback Electronics Letters V 12 #11 (1976)

Moorer, J. About this Reverberation Business Computer Music Journal V3 #2 p13-28 (1979) reprinted in Foundations of Computer Music ed. Curtis Roads and John Strawn MIT Press 1985 pp 605-639


Stautner, J.; Puckette, M.: Designing Multi-Channel Reverberators Computer Music Journal V6 #1 (1982) p 52

Law K.A. -masters thesis in EE Perdue University

Smith J.O. -waveguide reverberators -CCRMA Stanford can supply the whole works
Oppenheimer L. Digital Reverb -a Modern Miracle parts 1-5 Mix V 9 #5,6,7,8,9 (1985)

Strom S., Dahl H., Krokstad A., Eknes E.; Acoustic Design of the Grieg Memorial Hall in Bergen Applied Acoustics 18 1985 127-142

Lamoral R., Cremer L., Futterer T.; L' Acoustique du Palais Acropolis a Nice Acoustica V 86 pp 75-83 (1986)



Measured pulse response using pulses from a pulse generator with octave frequency 2000 Hz. The loudspeaker was positioned in the middle of the orchestra enclosure. The finer details in the echograms were investigated using a highly directional microphone. This microphone unit had a parabolic screen with a diameter of 1 m. The first strong reflection seen in the two echograms (time delay 30-35 msec) comes from the ceiling. The next strong reflections (time delay 45-120 msec) come from the side wall reflectors and the freely hanging reflectors below the ceiling. They therefore all have a lateral incidence.

Figure 1. Echogram of the Grieg Memorial Hall in Bergen, Norway (from Strom et al)

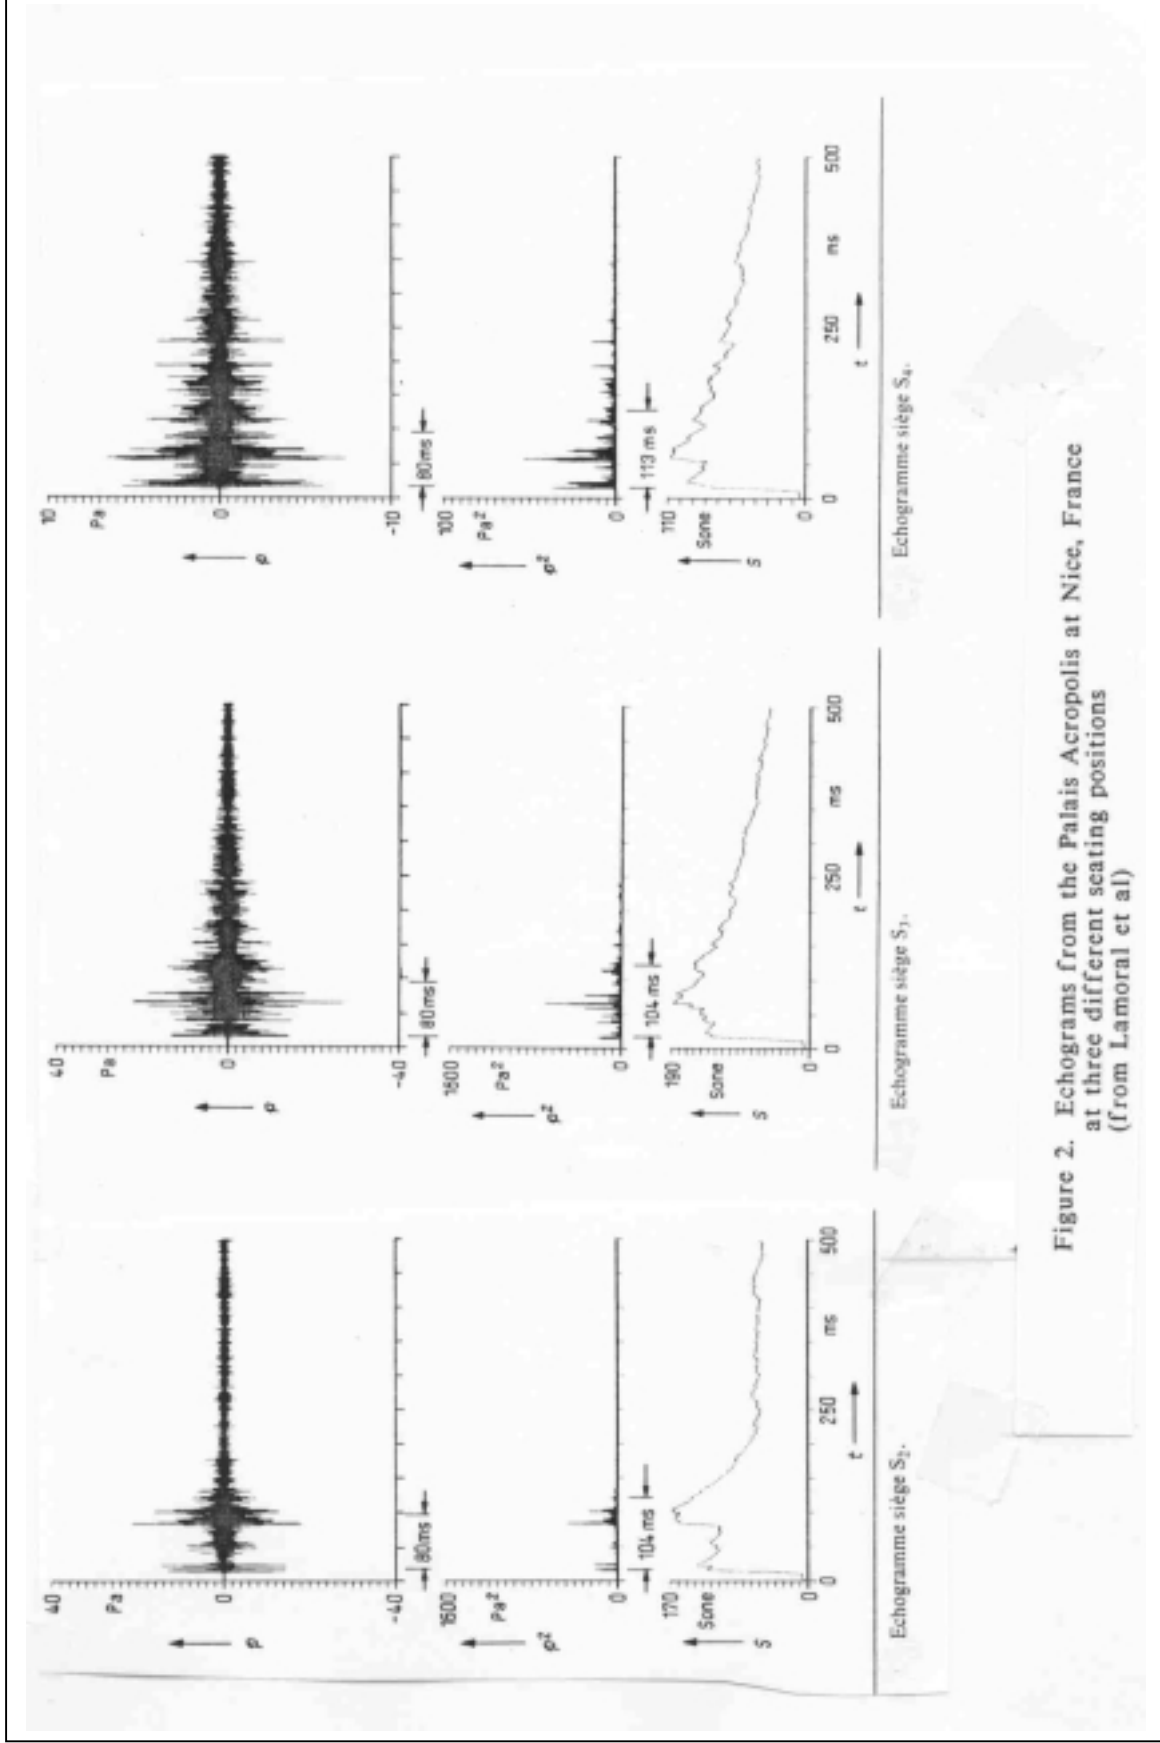


Figure 2. Echograms from the Palais Acropolis at Nice, France at three different seating positions (from Lamoral et al)



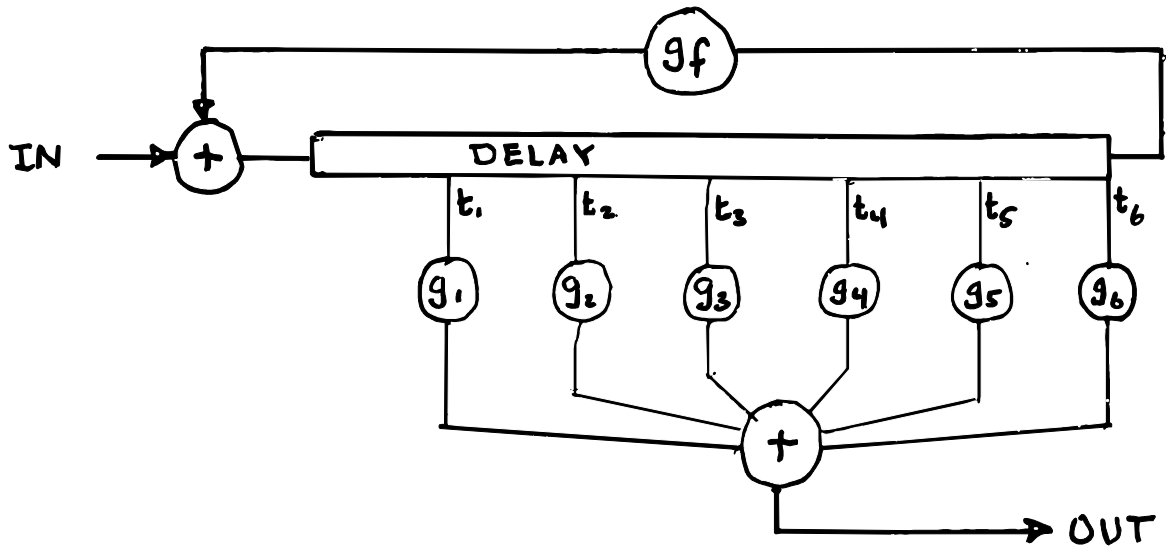


Figure 5. Delay with multiple feedback paths
 Note the density of the echos in the impulse response increases with time, but the amplitudes fluctuate (from Axon et al)



Impulse Response




Figure 4. Delay with a single feedback and multiple outputs
 If the output gains are chosen correctly, exponential decay results, but with a repeating pattern
 (from Axon et al)





Figure 5. Delay with multiple feedback paths
 Note the density of the echos in the impulse response increases with time, but the amplitudes fluctuate (from Axon et al)

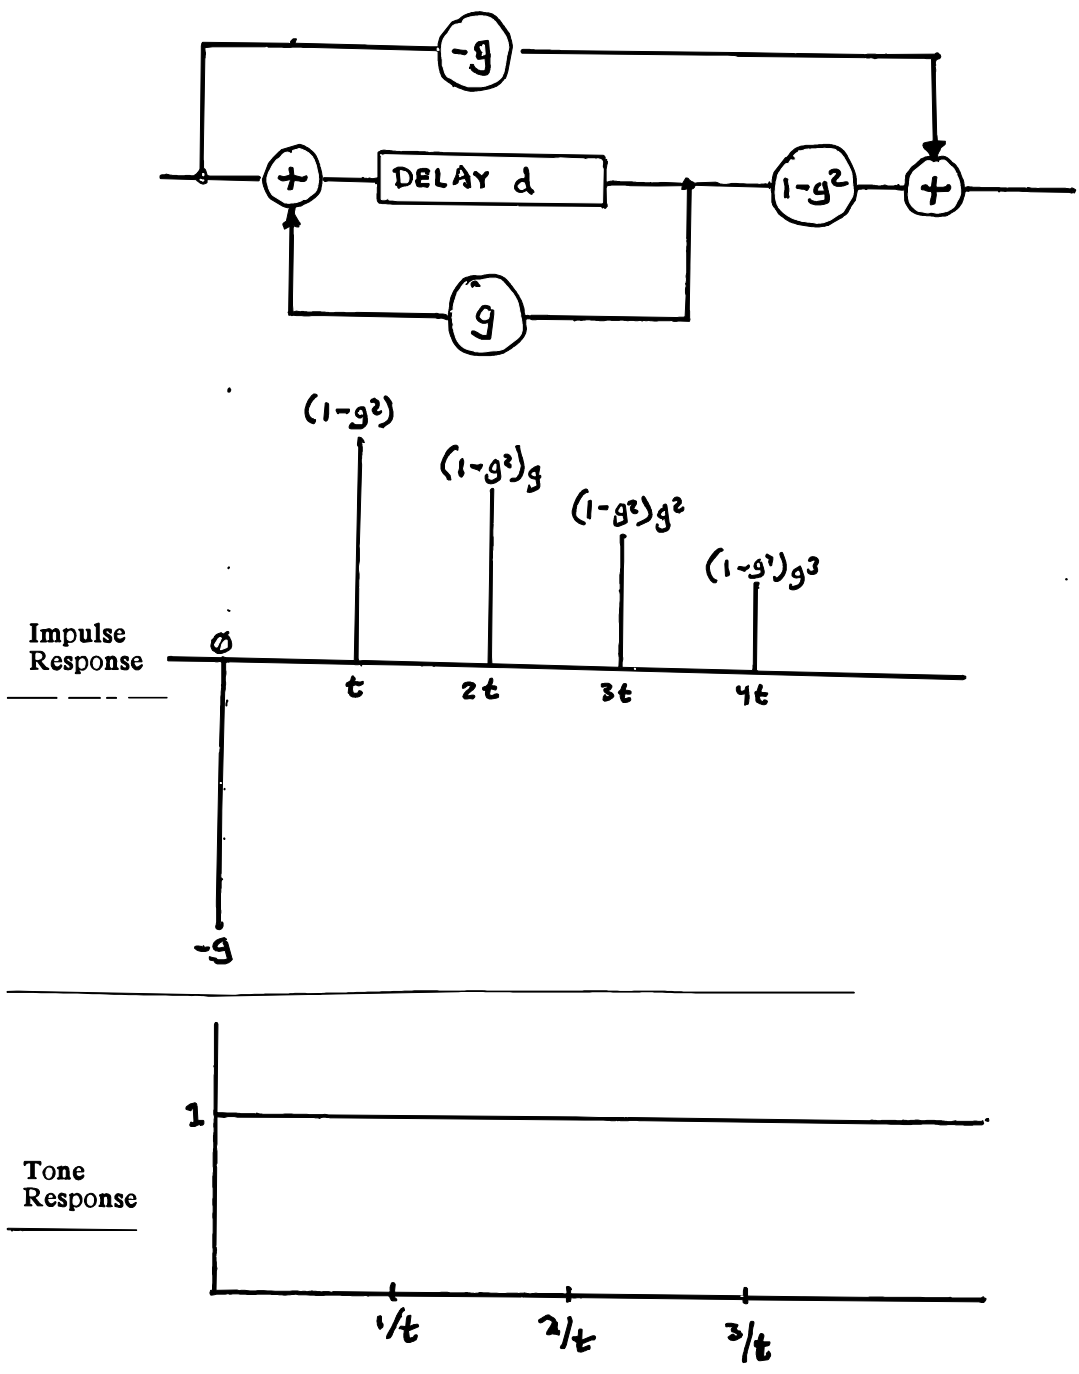


Figure 6. All-Pass network (from Schroeder)

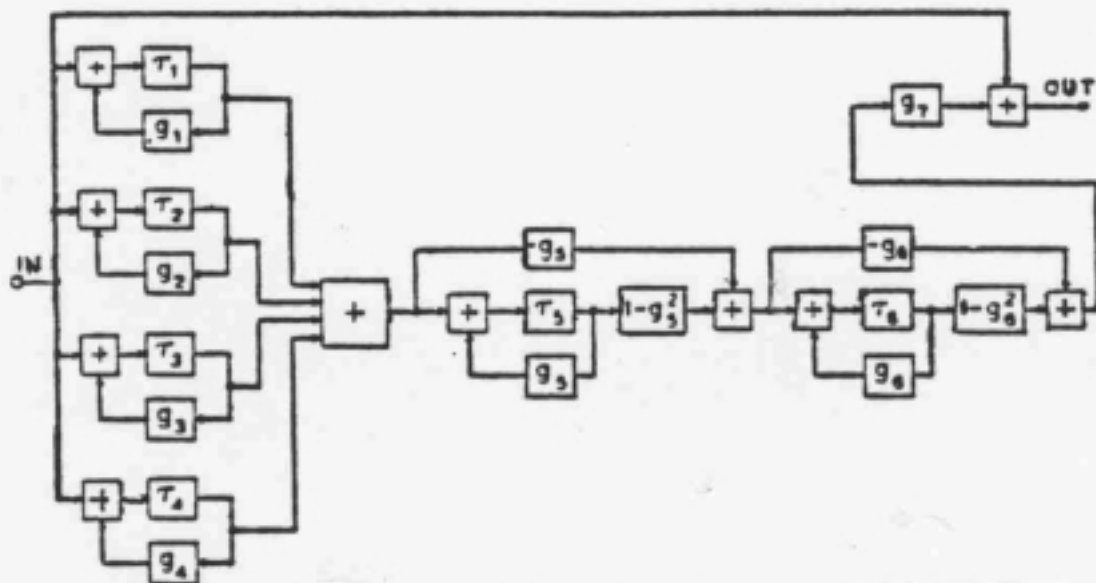


Figure 7. The Schroeder reverberator - two all-pass filters in series with 4 combs in parallel




Figure 8. Stautner's reverberator - four interconnected combs