

About This Reverberation Business

James A. Moorer
IRCAM
31, Rue St. Merri
75004 Paris, France

Introduction

When music is performed in a concert hall, a torrent of echos from the various different surfaces in the room strikes the ear, producing the impression of space to the listener. This effect can vary in evoked subjective response from great annoyance or even incomprehensibility, in the case of speech presented in a highly reverberant auditorium, to sheer ecstasy, in the case of late romantic music in the Vienna Grosser Musikvereinsaal. Most music is heard these days either in the comfort of one's home (or car), or in a university auditorium, both of which have generally short reverberation times. For this reason, most recordings of music already have some amount of reverberation added before distribution either by a natural process (*i.e.* recordings made in concert halls) or by artificial processes (plate or spring reverberators). Computer music is an especially fertile ground for artificial reverberation in that it is rarely, if ever, performed in highly reverberant concert halls. The use of the computer for the simulation of this reverberation also allows the composer another degree of freedom, namely to "tailor" the reverberation to the particular aural effect he wishes to present, allowing, for instance, each individual sound in the piece to carry an entirely different spatial aspect, if desired. In this paper, we review some of the work that has been done in the production of artificial reverberation by computer, and we present the fruits of our own labors along this line, both in the attempted simulation of the concert hall environment by computer and in the proliferation of different circuits for the realization of simulated room reverberation.

Historical Review

Although the study of acoustics may have started as early as the sixth century B.C. with Pythagoras' inquiries upon the discovery that halving the length of a string seemed to double its pitch, our own treatment will begin with the first published computer simulations of room reverberation by Schroeder (1961, 1962). We will not have time to discuss directly the first published study of room reverberation by

Sabine (published originally in 1900, reprinted in 1972). The two systems proposed by Schroeder are combinations of two unit reverberators, shown in canonical form in Figure 1. The first one is the all-pass filter, shown in Figure 1a in the one-multiply form. The second is the comb filter in Figure 1b.




Figure 1. The simplest unit reverberators: (a) one-multiply all-pass; (b) comb filter. The magnitude of g must be less than 1 for stability.

We should say a quick word about our use of signal flow graphs to represent filter structures. In our formulation, X always represents the input to a filter and Y always represents the output. An arc of a graph which is not labeled is a unity gain arc. Sometimes we insert extra nodes for supposed clarity. If more than one arc comes together at a node, they are to be summed. When several arcs leave a node, the same signal is supplied to all of the arcs. An arc can be a gain (a multiply), a delay (denoted by Z , the unit advance operator, raised to a negative power), or another filter, usually represented by a capital letter as a function of z .

To continue with our exposition, the unit reverberators shown in Figure 1 have, by themselves, the following transfer functions:

$$T(z) = \frac{g + z^{-m}}{1 + gz^{-m}} \quad (1)$$

$$T(z) = \frac{z^{-m}}{1 - gz^{-m}} \quad (2)$$

It is simple to show that the transfer function in Equation (1) is indeed that of an all-pass filter, in that the coefficients in the numerator are in reverse order of those in the denominator, thus forcing the zeros to be the reciprocals of the poles. This is sufficient for the filter to be an all-pass. We must remember, however, that the all-pass nature is more of a theoretical nature than a perceptual one. We should not assume, simply because the frequency response is absolutely uniform, that the filter is perceptually transparent. In fact, the phase response of the all-pass filter can be quite complex. The all-pass nature only implies that in the long run, with steady-state sounds, the spectral balance will not be changed. This implies nothing of the sort in the short-term, transient regions. In fact, both the all-pass and the comb filter have very definite and distinct "sounds" that to the experienced ear can be immediately recognised.

To use these as a reverberation generator, we must combine them in some manner. Schroeder suggested two different combinations, shown in Figures 2a and 2b. In these figures, A_i represents an all-pass filter and C_i represents a comb filter. Figure 2a thus shows a series connection of all-pass filters with a final addition of some proportion of the original signal. With the all-pass, there is also a feed-through of the original signal scaled by the gain g of the all-pass, so with the circuit of Figure 2a we will find a total direct signal contribution of k plus the product of the gains of the individual all-passes. In Figure 2b, we see a parallel connection of four comb filters followed by two all-passes, again with a feed-forward connection of a proportion of the original signal. The idea with both of these is to use the unit reverberators to simulate the effect of wall reflections and transit time of the wave front as it passes from one wall to the other. The addition of the direct signal simulates the proximity of the source to the destination. As the destination listener moves away from the sound source, the perceived reverberation remains at about the same amplitude, but the direct sound intensity decreases by the usual reciprocal distance squared term; thus there is a distance at which the direct and reverberant sounds are equal in amplitude, such that at larger distances than this, the reverberant sound field dominates. In these circuits, we attempt to simulate the wall reflections with the feedback paths and the transit time between reflections with the delay lengths. We shall see later that this is a somewhat crude way to simulate actual rooms. Both Schroeder and this author have proposed refinements (with corresponding increases in computer time required).

In any case, the use of the reverberators in Figures 2a and 2b implies that the various parameters (gains and delay lengths) of all the unit reverberators must be chosen somehow.




Figure 2. Combinations of unit reverberators to make a functioning reverberator. In each combination, the direct signal is added in with a scaling of k . (a) Series connection of all-pass networks. The functional dependence on z has been dropped in the illustration. (b) Parallel combination of comb filters followed by series combination of all-pass filters.

It is generally thought important to make all the delay lengths mutually prime, in that this reduces the effect of many peaks piling up on the same sample, thus leading to a more dense and uniform decay.

More recently, Schroeder (1970) reported on a way of simulating room reverberation while taking into account the exact geometry of the room. This study involves the simulation of the source as an omnidirectional pulsating circle (the study was done in two dimensions), following the paths of some 300 individual rays that are cast at equally spaced angles about the source. The presumptive linearity of the air and walls makes it possible to transmit a single ideal impulse, thus obtaining the impulse response of the simulated room. The impulse response can then be convolved with musical sound to produce the sound of that music in that room. Following this, he suggested that a new reverberator could be made that simulates exactly the first few echos, then simulates




Figure 3. Form suggested by Schroeder (1970) for simulation of early echos by direct summation, using a standard reverberator as shown in figure 2a or 2b for the late echos.

the later echos by means of a reverberator such as has been shown above. This design is shown in Figure 3. The first delays and gains are chosen following the geometry of the room; then the reverberator $R(z)$ is chosen to simulate the decay of the room reverberation after the density has reached the level where the individual pulses cannot be separated. We will come back to this idea later.

Some New Unit Reverberators

In the course of studying the art of making artificial reverberation, we have explored several new unit generators. These include the oscillatory all-pass in Figure 4a and the oscillatory comb in Figure 4b. In addition, we may also replace the gain terms with filters to produce decays, the lengths of which depend upon the frequency. For the non-oscillatory case, these are shown in general form in Figure 4d for the all-pass and in Figure 4c for the comb. There exist straightforward extensions to the oscillatory cases of comb and allpass for these configurations also. What must be discussed now is what the general utility of these units is and when they should be used. In fact, we have found that the only really useful one is the comb form in Figure 4c, but we include a discussion of the others for completeness and additional insight as to the problems involved. These reverberators are all modifications of Schroeder's original designs (Schroeder, 1961, 1962), and can all be described as various forms of recirculating delays, since they all use delays with feedback and sometimes feedforward. We should point out that these units were known as early as 1966 when H. Date and Y. Tozuka simulated these units using analog filters and variable-speed multi-head tape recorder delay. Needless to say, the filter realizations were somewhat complex.

The reverberators in Figures 2a and 2b have been used extensively at Stanford University, at IRCAM, and many other sites around the world in varying different configurations. With the all-pass reverberator of Figure 2a, several problems were noticed by this author along with L. Rush, D.G. Loy, K. Shoemaker, J. Chowning and others:

- 1) The decay did not start with a dense sound and die out slowly in an exponential manner. In fact, the higher the order, the longer it took for the density to build up to a pleasing level. This produces the effect of a lag in the reverberation, as if the reverberation followed the sound by some hundreds of milliseconds.
- 2) The smoothness of the decay seemed to be critically dependent on the choice of the parameters involved: the gains and delay lengths of the individual unit reverberators. Just changing one of the delay lengths from its prime number length to the next larger prime number, which could be a change as small as 2 samples, has been noticed to occasionally make the difference between a smooth-sounding decay and a ragged-sounding decay.
- 3) The tail of the decay showed an annoying ringing, typically related to the frequencies implied by some of the delays. This produced a somewhat metallic sound that was generally found objectionable.

It would be nice if we could somehow derive these subjective properties directly from the transfer functions of the filters. In this way, we would be able to lift the evaluation of these




Figure 4. Some new unit reverberators. (a) is an oscillatory all-pass. (b) is an oscillatory comb. The impulse response of the oscillatory units is a train of pulses the amplitudes of which correspond to a decaying sinusoid. (c) is a comb filter with another filter inside the loop, identified by its transfer function $T(z)$. (d) is an all-pass with filters in the loops identified by their transfer functions: $H(z)$ in the feed-forward path, and $S(z)$ in the feed-back path. For this to remain all-pass, these filters must be complex conjugates of each other.

out of the domain of empirical study and to design entirely on paper what should be reverberators which sound good. Unfortunately, there does not seem to be any viable way to do this at this time. Our perception of the quality of

the reverberator does not seem to be simply related to the quantities that one usually derives from transfer functions, such as group delay or root locus.

Although the reverberator of Figure 2b did not show exactly these same problems, it shows other problems, namely: 1) Any attempt to reverberate very short, impulsive sounds, such as drum strokes, gave distinct patterns of echos rather than smooth reverberant sounds. The sound was somewhat like flutter echo in rooms with parallel walls. 2) The decay continued to show a metallic sound, especially with longer reverberation times.

Our first thought to help this situation was to try the oscillatory networks shown in Figures 4a and 4b. The transfer functions of these two unit reverberators are then as follows:

$$H(z) = \frac{g_1 + g_2(1 + g_1)z^{-m} + z^{-2m}}{1 + g_2(1 + g_1)z^{-m} + g_1z^{-2m}} \quad (3)$$

$$H(z) = \frac{z^{-2m}}{1 - g_1z^{-m} - g_2z^{-2m}} \quad (4)$$

Both these units can be described completely by the decay rate of the exponential envelope and the frequency of the oscillation. The decay rate r is the factor by which the amplitude is diminished every m samples (in the impulse response). The frequency θ is that of the sinusoidal part of the decay. In the case of the oscillatory comb, we can compute the coefficients g_1 and g_2 from the frequency and decay rate as follows:

$$g_1 = 2r \cos\left(-\frac{\theta}{m}\right) \quad (5)$$

$$g_2 = -r^2 \quad (6)$$

Negative decay rates can also be used, making each pulse of the impulse response alternate in sign. The condition for stability is that the magnitude of r must be less than one. We can now design with the oscillatory comb just the way we have with the simple comb, with r taking the logical place of g , and with one more parameter to choose. The frequency θ is divided by m in the equation because of the delay in m samples between each pulse in the impulse response.

We can likewise compute the coefficients of the oscillatory all-pass from the desired decay rate and frequency as follows:

$$g_1 = r^2 \quad (7)$$

$$g_2 = -2 \frac{r}{1 + r^2} \cos\left(\frac{\theta}{m}\right) \quad (8)$$

In the case of the all-all-pass reverberator (or series-all-pass, Ed.), such as is shown in Figure 2a, the use of the oscillatory all-pass made little or no perceptible difference, even when all units were so modified. There was a slight difference in the tail end of the decay in that the dominance of one pitch over another kept shifting so that we heard the pitch of one delay for a short while, then the pitch of another delay, and so on.

With the combination comb-all-pass reverberator, such as is shown in Figure 2b, the use of oscillatory combs did help the decay a bit in the same way as with the all-all-pass case, but did not correct the principal problem, which was the response to short, impulsive sounds. The improvement is enough, however, to say that short of placing a filter in the loop, which will be discussed subsequently, the comb-all-pass reverberator with oscillatory combs was the best-sounding reverberator up to this point.

Figures 4c and 4d show unit reverberators with filters in the loops. Figure 4c shows a comb-type, and Figure 4d can be made an all-pass if the loop filters $H(z)$ and $S(z)$ are complex conjugates of each other. This implies that their impulse responses are the time reverse of each other. Since the feed-forward filter $H(z)$ passes directly into the output, it should be unconditionally stable. This gives us the entertaining result that either $H(z)$ and $S(z)$ must be FIR filters, or $S(z)$ must be an unstable filter when viewed in isolation. Strangely enough, if the magnitude of $S(z)$ when evaluated on the unit circle is everywhere less than unity, the unit reverberator as a whole will be stable (see below). This notwithstanding, we consider it a somewhat bad idea to

Table 1. Dependence of absorption coefficients on humidity and frequency.

Relative Humidity	Frequency in Hertz			
	1000	2000	3000	4000
40	.0013	.0037	.0069	.0242
50	.0013	.0027	.0060	.0207
60	.0013	.0027	.0055	.0169
70	.0013	.0027	.0050	.0145

(Above) Attenuation Constant m of air at 20° C and normal atmospheric pressure in units of meter⁻¹. (Below) The same table in units of dB per meter.

Relative Humidity	Frequency in Hertz			
	1000	2000	3000	4000
40	.0056	.0161	.0300	.1051
50	.0056	.0117	.0261	.0899
60	.0056	.0117	.0239	.0734
70	.0056	.0117	.0217	.0630

put an unstable filter in the loop because the delay of m samples prevents any annihilation of the response of $S(z)$ for a certain time, during which the response of $S(z)$ may well have exceeded the bounds of the computer on which it is being run. For this reason, we consider it more reasonable to use FIR filters for $H(z)$ and $S(z)$. To make the circuit realizable, we must insert a delay after the node where the feedforward path breaks from the original signal and before the node where the feedback path adds into the original signal, a delay equal to the total delay of the feedforward filter $H(z)$. Some additional efficiency can be obtained by the use of linear phase filters here.

The purpose of placing a filter in the loop is to simulate the effect of the attenuation of the higher frequencies by the air. This attenuation is due "partly to the effects of viscosity and heat conduction in the gas and partly from the effects of molecular absorption and dispersion in polyatomic gases involving an exchange of translational and vibrational energy between colliding molecules" (this description lifted with minor paraphrasing from Beranek (1954)). Unfortunately, we have found much disagreement in the literature as to exactly what these constants are. The reported values vary by as much as a factor of five in the decay rate. In any case, the intensity of the sound is generally thought to vary, for any given frequency, according to the following formula:

$$I(x) = \frac{1}{x^2} I_0 e^{-mx}$$

where I_0 is the intensity at the source, x is the distance from the source in meters, and m is the attenuation constant which depends on frequency, humidity, pressure, and temperature. Table 1 was taken from Kuttruff (1973) who claims to have taken it from Harris (1963) and Harris and Tempest (1964). The absorption seems to go up even higher at lower values of humidity and at higher frequencies. These absorption coefficients may not seem very high, but in a large concert hall such as the Boston Symphony Hall, the reverberation time of more than 1.7 seconds implies that the sound has traveled more than 600 meters in traveling between the various walls. This would mean (at 40 percent relative humidity, a typical figure in the Boston area during the concert season) that the 4 KHz signals would have been attenuated upwards of 60 dB more than the 1 KHz signals. This just shows that this process is significant and cannot help but have a substantial effect on the resulting spectrum of the late decay.

Now we can approach the question of modifying the unit reverberators to simulate this facet of sound transmission. We must keep in mind, however, certain perceptual relations in reverberation that have been noticed over the years. During the performance of a passage of music, only the first 10 to 20 dB of the decay of the reverberation is typically audible because of the changing nature of the music. It is only in pauses or at ends of movements that the entire decay of the reverberant sound will be audible. It is not clear what this fact implies except perhaps that we must concentrate equally on the early decay and the late decay for a comprehensive simulation.

Since one problem that arises immediately with the new unit generators is that of the stability of the resulting unit

when one inserts a filter into the feedback loop, it is important to take a closer look at the stability conditions for such filters. We can compute the stability of the comb by carrying the transfer function of the filter that has been inserted in the loop along explicitly, then separating it into its magnitude and phase functions as follows:

$$G(z) = \frac{z^{-m}}{1 - gT(z)z^{-m}} \quad (9)$$

This then has as magnitude function, evaluated on the unit circle:

$$G(e^{j\omega P_s}) = \frac{1}{1 + g^2 |T|^2 - 2g |T| \cos(\theta(\omega) - m\omega P_s)} \quad (10)$$

where $|T|$ is the magnitude of $T(e^{j\omega P_s})$ and $\theta(\omega)$ is the angle of $T(e^{j\omega P_s})$. ω is the radian frequency around the unit circle and P_s the sampling period in seconds.

Ignoring for a moment the angular term, we can see that $|T|$ follows g around at each point in the formula. This implies that we can consider the effect of the filter separately on the radii of the roots and on their angular placement. Without the filter, the roots are placed evenly at multiples of the frequency $(1/m)P_s$ and at a uniform radius of the m th root of g . With the filter, the radii have been changed to $g|T|$, which depends on the input frequency ω , and the angular position has been distorted to the following:

$$\omega_i = \frac{2\pi i + \theta(\omega)}{mP_s} \quad \text{for } i = 0, 1, 2, \dots \quad (11)$$

For stability, all we need to do is assure that the radius never reaches 1.0, so the condition for stability becomes this:

$$g < \frac{1}{\max_{0 \leq \omega \leq 2\pi} |T|} \quad (12)$$

Thus knowing the absolute maximum value of $|T|$, we can easily scale g to be unconditionally less than the reciprocal of this. This guarantees stability.

To test this idea, the filter shown in Figure 5 was incorporated into a reverberator of the form shown in Figure 2b. The loop filter is a simple first-order filter. For stability, the magnitude of g_1 should be less than one. If we insist that the loop filter be strictly low-pass, then g_1 should be positive, and




Figure 5. One form of the comb with a filter in the loop as shown in Figure 4c. Here we have chosen a simple one-pole filter. If g_1 is positive, it will be a low-pass filter. For stability, $g_2 / (1 - g_1)$ must be of magnitude less than one. (Note: if g_1 and g_2 are positive, this reduces to the condition that $g_1 + g_2 < 1$.)

the maximum of $|T|$ will occur at $\omega = 0$, as is shown by its transfer function:

$$T(z) = \frac{1}{1 - g_1 z^{-1}} \quad (13)$$

and its maximum value at $\omega = 0$ will be:

$$T(1) = \frac{1}{1 - g_1} \quad (14)$$

Thus if we set g_2 to $g(1 - g_1)$, where g is now between zero and one as before, the resulting filter is unconditionally stable and has the side benefit of giving us a parameter in the familiar zero to one range.

Again, the idea of inserting such a filter in the loop is to simulate the absorption of the high frequencies by the air. The frequency response of such a simple low-pass filter as that shown in Figure 5 can never conform to the spectral modification necessary to exactly simulate this property, but was chosen as a compromise with efficiency, since it only adds one multiplication. There is a question, however, of how to set this parameter g_1 which controls the roll-off of the filter. To this end, a series of optimizations was done using the Marquardt algorithm (Marquardt, 1963) to match the frequency response of this low-pass filter in a least-magnitude-squared manner to the actual data reported in the literature. The results are summarized in Table 2. The resulting optimal values of g_1 are shown in Figure 6. Unfortunately the values are strongly dependent on sampling frequency and somewhat dependent on humidity, temperature and pressure. The figures are given here for sampling rates of 10 KHz, 25 KHz, and 50 KHz at four different values of humidity. Again, the weakness of the first-order filter as a low-pass means that these can only be considered very vague approximations of the actual sound absorption of air, so the exact value of this coefficient should not be taken too seriously. These should only be considered guideline values.

Following these guidelines, we can choose a value for the low-pass filter coefficient that represents the attenuation that a specific delay will produce. Whatever delay that we choose for this comb, we can convert that delay into meters by multiplying by the speed of sound in air, which is generally taken to be 344.8 meters per second at 22 degrees Centigrade and standard pressure of 751 millimeters of mercury. For the appropriate value of humidity and sampling rate we can read off directly the value of g_1 from the graphs of Figure 6. For intermediate sampling rates, an interpolation can be done.

Following this procedure does in fact produce a reverberation such that the decay of the higher frequencies is somewhat faster than that of the lower frequencies, and the resulting sound is indeed somewhat more realistic. Several other benefits immediately occur also, such as the loss of sensitivity to the exact delay length and a more robust treatment of short, impulsive sounds. The configuration we have been favoring is that with six comb filters, each with low-pass filters in the loop, and a single all-pass filter.

Beranek (1962) reports that for "good" bass quality in a concert hall, it should have a ratio of the reverberation time at 125 Hz and the mid-frequency reverberation time (average of the 500 and 1000 Hz reverberation times) of somewhat greater than unity. The comb of Figure 5 exhibits exactly this property for any positive value of g_1 . Specifically, for the values of g_1 in Figure 6, the ratio is around 1.02, which is designated by Beranek as customary in a "good" concert hall. This point is actually somewhat moot in the computer music domain, in that the reason that bass is to be emphasized in concert halls is that string basses are somewhat weak, especially when one considers the lack of sensitivity of the ear to lower frequencies, such that they must be amplified by the room itself if possible to balance properly with the orchestra.

We have found that in the use of the six-comb reverberator, certain distributions of values make more sense than others. For instance, with more than six combs or more than one all-pass, the difference in sound quality is largely unnoticeable. The delays seem to work well when distributed linearly over a ratio of 1:1.5 with all the g_2 terms set to a constant number g times $(1 - g_1)$. The all-pass seems to work well at about a 6 millisecond delay and a gain of 0.7 or so. As a concrete example, Table 2 shows the values of delays and low-pass filter gains for such a reverberator at 25 KHz samp-

Table 2. Examples of parameters for actual reverberators of the type shown in Figure 2b, using six comb filters with the low-pass filters in the loop, such as shown in Figure 5, and one standard all-pass following, as shown in Figure 1a. In each of these cases, the all-pass should be set to a 6 millisecond delay with a gain of 0.7 or so.

		25 KHz	50 KHz
Delay (in ms.)		g_1	g_1
COMB 1	50	0.24	0.46
COMB 2	56	0.26	0.48
COMB 3	61	0.28	0.50
COMB 4	68	0.29	0.52
COMB 5	72	0.30	0.53
COMB 6	78	0.32	0.55

ling rate, and at 50 KHz. Again, the loop gains g_2 should be set to a number g times $(1 - g_1)$ for each comb. Each comb will have a different value of g_1 , but should have all the same value for g . This number g determines the overall reverberation time. For example, values around 0.83 seem to give a reverberation time of about 2.0 seconds with these delays.

The distribution of these numbers g for the combs determines the dominance of one comb over another in the decay. If one notices that a particular comb's pitch is dominating the decay, the value of g for that comb can be reduced. With less than six combs, one can never arrive at a set of values of g that mask entirely the dominance of a particular comb.


Additionally, the all-pass seems to be a somewhat sensitive subject too, in that too short a delay for the all-pass gives an annoying "puff-puff" response to background noise. If there are any clicks or pops in the sound being reverberated, a short all-pass will surround each click with a "puff" consisting of its own impulse response, sounding not unlike a very quiet cymbal crash. To eliminate this kind of sensitivity, shorter delays should be avoided. A further complication, however, is that any delay longer than six milliseconds seems to produce an audible repetition period; thus, the delay of the all-pass is pretty well limited to six milliseconds.

Turning now to the reverberator as a whole, we have tested this formulation with delays as short as 10 to 15 milliseconds for the combs with the result that the sound is still good, and that long reverberation times can still be obtained even with such short delays. The subjective impression of the simulated concert hall changes in this case such that at the shorter delay values, one might well imagine that one were inside a garbage can rather than the Symphony Hall, but the density and "naturalness" are still good, although one might quarrel with the subjective impression of the timbre.

Indeed, this seems to be about as good as one can do with the recirculating delay system of reverberation. It does not seem to add anything by putting more complicated filters in the loop, except that one can then simulate other architectural decisions, such as the compositions of the walls or the absorption of an audience. The low-pass filter in the loop seems to help the response to short, impulsive sounds by smearing out the echos. Each echo is no longer a discrete impulse, but is lengthened by the simulated transmission characteristics of the air itself, so that even in the early reverberation period with its correspondingly lower density, the discrete echos seem to be effectively masked. It likewise did not seem to add anything to use the all-pass network with filters in the loop (Figure 4d), so this unit does not appear useful.

Figure 6. Graphs of the optimal gain, g_1 , of the low pass filter in the feedback loop of a comb filter as shown in Figure 5 for simulating as well as possible the natural high-frequency attenuation of the air. Figure 6a is for a sampling rate of 10 KHz., Figure 6b, for 25 KHz., and Figure 6c, for 50 KHz. The delay length must be converted to meters, then for the appropriate humidity (higher on the east coast), the value of the coefficient can be read off directly. For intermediate sampling rates or humidities, an interpolation must be done. Since this value is not highly critical, no special care must be taken with such an interpolation.

We might mention here that it is not absolutely necessary to set the values of g_1 from the optimal values shown in Figure 6. One can just set the values arbitrarily to produce different kinds of sounds that have little or no relation to physical reality. Larger values of g_1 produce a reverberation with a very bright beginning and a very muffled decay. Small



values of g_1 tend toward the case where there is no filter in the loop. The use of different values of g_1 for different combs produces no particularly different sound, but can cause dominance of one comb over the others.

We should add here a quick word about implementation. Normally, these reverberation schemes will be implemented either on special-purpose digital hardware or on general-purpose computers. In either case, one should pay special attention to the error characteristics of these reverberators, as this is important with all types of filters. Following the formulation of Oppenheim (Oppenheim and Weinstein, 1972; Oppenheim and Shafer, 1975), we can approximate the error amplification of the reverberator of Figure 5 as $1/(1-g)$. Likewise, the signal itself can be amplified by that same factor. For example, if $g = .75$, then we can expect the quantization error of the input signal to be increased by about a factor of four. This would imply that we should keep two or more extra low-order bits in the computation so that this error would be truncated at the output. Likewise, we would have to keep two extra high-order bits to assure that the total dynamic range did not exceed the word length of the machine. This implies that, for instance, to get 12 good bits of output, one would have to do the intermediate computations with 14 to 16 bits, or more. Likewise, to get 16 good bits of output, one would need even more bits internally. This is improved a little when we consider that generally the reverberated signal will be attenuated before it is added into the output, but in any case, one cannot afford to ignore the effects of finite register length in the implementation.

If this were not enough, there is a further phenomenon that occurs with fixed-point implementation of digital filters: that of limit cycles. Depending on the type of arithmetic used (rounding, truncation, etc.), the output of the filter (reverberator) may not descend exactly to zero even after the input has vanished. Output values of 1 in the last bit, or sequences of +1 and -1, are quite common in fixed point. This last case of alternating sign is especially annoying, because in 2's complement arithmetic, truncating the low-order bits does not cure the problem. Again, it is not a particularly disastrous effect, but one should not be surprised when it happens.






Figure 7. Diagram of a sound ray going from a sound source to a destination listener via a wall reflection. Just as in the case of light, we can think of the reflected sound as originating from a "phantom" source which is behind the wall, located on a line from the source that is perpendicular to the wall.



(a)




(b)


Figure 8. Patterns of phantom sources for two room shapes. (a) First and second bounce phantom sources with the perpendiculars of construction shown. (b) Pattern of phantoms up to the fifth bounce. Note that even though there are no parallel walls, the phantom sources show very regular patterns.

But What About Real Rooms?

The only problem with these systems of recirculating delays is that they can never correspond to the room reflection pattern in a real room. This is because the reflections in a room, even a square one, do not typically come in regular sequences separated by equal amounts of time. To see why this should be so, we have to take a brief look at the science of room acoustics. Since we do not have room here to expound at length on the subject (the interested and brave



(a)



(b)

Figure 9. The phantom sources in a corner fall in a circle the center of which is the vertex of the corner. In (a) we can trace five phantoms which start with the first reflection off the upper wall, which are labeled 1a, 2a, 3a, 4a, and 5a. Likewise, there are five phantoms starting with an initial reflection off the bottom wall, which are labeled 1b, 2b, 3b, 4b, and 5b. (b) shows a rectangular corner, wherein the phantom sources number only three, one of which is a superposition of the two-bounce sources from each wall.

reader is referred to one of the many introductory texts on the subject, such as Beranek (1954) and Kuttruff (1973)), we will only describe a small corner of what is called geometric acoustics.

This entire section is largely tutorial in nature and presents a somewhat naive view of acoustics that will subsequently be shown to be useless for simulating the sounds of real rooms. The more advanced or impatient reader might want to skip right away to the conclusions in the next section.

Much as in the study of light, we can keep track of the reflections in a manner that depends only on the position of

the source and not on the position of the destination listener. To compute the actual impulse response of the room, we must specify the position of the destination listener; but by using the method of phantom sources, we can keep track of the reflections independent of the position of the listener up until the last moment. To see how this works, let us turn to Figure 7, wherein we see a simplified case consisting of one wall, one source, and one destination. In the case shown, direct sound omitted, the wave reflects off the wall and strikes the destination. This is equivalent in effect to the case where the wall does not exist and there is another identical source that spoke at the same time as the true source. This "phantom" source is located on a line connecting the destination point and the reflection point. We may also locate the phantom source by drawing a line from the source that is perpendicular an equal distance on the other side of the wall. Thus if we now consider a room with three walls, shown in Figure 8a, we can trace the various phantom sources given only the position of the source and the room geometry. In doing this for a two-dimensional room, we are assuming that the floor and ceiling are infinitely absorbing and thus do not reflect. We see that the phantom sources are represented by their bounce number. Each first bounce phantom source reflects off the other two walls to make six "second-bounce" sources. Figure 8b shows an ensemble of phantom sources up to the fifth bounce.

To give some slight additional insight into how these phantom sources work, we can consider the case of a corner in isolation such as shown in Figure 9a. In this case, all the phantom sources will be located on a circle the radius of which is the distance from the source to the vertex. In a corner, there will only be a finite number of reflections. A phantom source cannot reflect off a wall which it is behind (obviously!). Thus, as we follow the reflections in Figure 9a starting with the reflection off the upper wall, we find phantom sources at positions 1a, 2a, 3a, 4a, and 5a, the number corresponding to the number of bounces. Likewise, starting with the lower wall, we find the phantom sources 1b, 2b, 3b, 4b, and 5b, after which the phantom source is behind the wall which it would reflect off of, thus terminating the sequence. We can also see, in Figure 9b, the well-known case of a right angle where there are exactly three phantom sources, of which two are first bounce sources and the third is a second bounce source which receives contributions from both the first bounce sources. With any angle greater than 90° , any given corner will create only two phantom sources.

We can easily write a computer program to keep track of these phantom sources. To do so, at least for rooms which are convex polygons, we only need to keep track of the locations of the phantom sources, and be able to determine whether the phantom source is in front of or behind a given wall when we are computing new phantom sources. To review our coordinate geometry a bit, let us recall that the normal form for a line in the Euclidian plane is the following:

$$Y \cos(\theta) - X \sin(\theta) = \rho \quad (15)$$

where θ is the inclination of the line and ρ is the perpendicular distance from the line to the origin. Given the endpoints of a

line as (X_0, Y_0) and (X_1, Y_1) , we can determine the equation for the line as follows:

$$\text{let } d = \sqrt{(Y_1 - Y_0)^2 + (X_1 - X_0)^2}$$

$$\sin(\theta) = \frac{Y_1 - Y_0}{d} \quad (16)$$

$$\cos(\theta) = \frac{X_1 - X_0}{d} \quad (17)$$

$$\begin{aligned} \rho &= Y_0 \cos(\theta) - X_0 \sin(\theta) \\ &= Y_1 \cos(\theta) - X_1 \sin(\theta) \\ &\equiv \frac{Y_0 X_1 - X_0 Y_1}{d} \end{aligned} \quad (18)$$

Now that the line is in normal form, we can test any given point (X_t, Y_t) to determine which side of the line it is on. That is,

$$Y_t \cos(\theta) - X_t \sin(\theta) \geq \rho$$

means that (X_t, Y_t) is on one side of the line, whereas

$$Y_t \cos(\theta) - X_t \sin(\theta) < \rho$$

means that (X_t, Y_t) is on the other.

We notice that with the normal equations for a line, there is always an ambiguity. We can always add π to the angle and negate ρ , giving a different equation for the same line. This is less of an ambiguity than a fortunate extra degree of freedom, for we can use this to orient the wall equations in a standard way such that for any point (X_t, Y_t) that is inside the polygon (again, assuming convexity), with N walls whose equations are determined by θ_i and ρ_i , the following will be true:

$$Y_t \cos(\theta_i) - X_t \sin(\theta_i) > \rho_i \quad \text{for } i = 1, 2, \dots, N \quad (19)$$

All we need to do to allow us to use this simple test is to go through the wall equations with any one point (X_d, Y_d) that is known to be inside the polygon and flip any equation for which the above does not hold. For a convex polygon, we can construct such an "internal" point by averaging the X and Y coordinates (separately, thank you) of all the vertices. Given the normal equations for the wall lines, we can now decide whether a given source will reflect off a given wall just by verifying that the above condition is satisfied for the appropriate value of i .

Continuing in this manner, we can compute the position of the new phantom source just by taking the equation of a line going through the originating source (phantom or not) that is perpendicular to the wall. If the equation of the wall is as shown in Equation (15), then such a line is given by the following equation:

$$Y \sin(\theta) + X \cos(\theta) = Y_s \sin(\theta) + X_s \cos(\theta) = \rho_s$$

where (X_s, Y_s) are the coordinates of the originating source. If we then intersect this with the wall line, we can find the coordinates of the intersection (X_x, Y_x) as the following:

$$X_x = -\rho \sin(\theta) + \rho_s \cos(\theta)$$

$$Y_x = \rho \cos(\theta) + \rho_s \sin(\theta)$$

Thus the differences in the coordinates are simply:

$$\Delta X = X_x - X_s$$

$$\Delta Y = Y_x - Y_s$$

and the coordinates of the new phantom source on the other side of the wall are just the following:

$$X_n = X_x + \Delta X = -2\rho \sin(\theta) + Y_s \sin(2\theta) + X_s \cos(2\theta)$$

$$Y_n = Y_x + \Delta Y = 2\rho \cos(\theta) - Y_s \cos(2\theta) + X_s \sin(2\theta)$$

Now that we know how to keep track of the sources, we must see how this affects the listener. There are three main sources of attenuation of the sound from a phantom source. The first is that each bounce presumably reduces the amplitude a bit. Wall reflection characteristics are actually very complex, depending on both the frequency and the angle of incidence of the wave front. We will brush aside this nasty intrusion of nature and for the time being represent the attenuation of each wall by a single number, the reflection coefficient. Another source of attenuation is just the fact that any spherical wave is diminished over some distance at a rate proportional to the reciprocal of the distance it has travelled. (Note that we are working in amplitude here, and not intensity. For intensity, it is the reciprocal of the distance squared.) The last source of attenuation is that of the air itself, which we have already discussed in the previous section.

If we ignore for a moment the frequency dependence of everything, then we see that it should be quite straightforward to compute the contribution to the listener from each phantom source. We must only assure ourselves that the ray from the phantom source to the listener crosses the wall that produced the phantom source. Otherwise, the listener

would not be able to “see” that phantom source. Thus at each step we can compute the effect at the listener by adding up the contributions of all the phantom sources.

A program was written here at IRCAM for examining two-dimensional convex polygonal rooms using this simplified phantom source model. The procedure was to first construct the impulse response of the room and then later to convolve this directly with the desired signal to produce the effect of that sound in this fictitious room. When we started making trials with this program, we found that the pattern of echos very quickly got extremely dense, so that after a certain time, every sample in the computed impulse response was non-zero. Visually, it greatly resembled white noise in the decay. One of many surprises was that although the strength of each phantom source diminished exponentially with each bounce, the overall contribution of all the phantom sources did not diminish nearly as rapidly because of the fact that the number of phantom sources grows exponentially also, thus helping to counteract the decreasing strength. Thus even though the strength of a given source was down by 120 dB, it might not be negligible because there happen to be 10,000 of them of roughly that same strength and same distance, thus raising the total contribution to only 40 dB down. This implies that the results of ray simulations of room reverberation (Schroeder 1970) should be carefully examined since they cannot take this effect into account. With ray simulations, one arrives at a total reverberation time that can be somewhat shorter than those calculated using the phantom source model we have described here.

Having written this program, we are now in a position to use it to gain some insight into how reverberation works. The first case we will discuss, however, is the classical example that needs no program: the rectangular room. This is one case where so many of the phantom sources fall on top of each other that they form a rectangular array, as is shown in Figure 10. This is similar to the barbershop mirror phenomenon, where one sees an endless array of phantom sources stretching off to infinity in both directions, some facing forward and some facing backward. The point of discussing this case, however, is to show that no artificial reverberator based on recirculating delays can ever be identical to even this geometrically simple case. This is immediately clear from putting a listener into the room. The contributions from just the row that crosses the room are at distances that are related by the square root of the sum of the squares of the vertical and horizontal distances. If it is granted that the vertical distance increases in regular increments of multiples of the room height, the resulting distances are not integer multiples of anything. In the limit as the sources get further and further away, the distances become almost equally spaced, implying that although the early response could never be simulated exactly by recirculating delays, it may be possible to approximate the late response. This also gives us some insight as to why typical artificial reverberation is always easily distinguishable from natural reverberation. We should take care to note that the case of the rectangular room is a very special one and we choose it as an example here just to show that the echo pattern will be different from anything we can produce using any of the usual combinations of recirculating delays.

Following an idea reported by Schroeder (1970), we may directly simulate the early response by a finite impulse




Figure 10. Phantom source patterns for a rectangular room form perfectly regular patterns of rows and columns. Any room shape which tessellates a surface by reflecting across each of its edges will have such a repeating pattern of sources. Such figures include equilateral triangles, for instance, and regular hexagons.

response filter, and simulate the late response with a standard reverberator such as is shown in Figure 3. Here we choose to simulate the first n delays separately with separate gains, then to send the delayed signal into a standard reverberator $R(z)$ with its own gain of g . We must then ask what are the design requirements of this reverberator and how do they differ from stand-alone reverberators. Our experience is, sadly enough, that they do not differ from stand-alone reverberators, in that they need all the usual complexity. Again, the most pleasing and natural-sounding system we have tried is the comb-all-pass combination with six comb filters, with low-pass in the feedback loops, and one all-pass. It is not necessary, however, to use the very long delays (for instance, the ten to fifteen millisecond range). All the design guides described in previous sections apply for these reverberators, including considerations of total reverberation time, ratio of minimum delay to maximum, etc., etc.

The rationale for simulating just the early echos is based on the various perceptual phenomena that ensue with the spacing of echos. With the case of two impulses spaced closely in time, the separation of these two impulses determines a wide range of perceptual effects. Certainly if the two pulses are more than 30 to 50 milliseconds apart, they will be heard as two separate and distinct pulses. Between this figure and one or two milliseconds, they will not be heard as distinct pulses, but the timbre of the combination will change (Green 1971). At less than one millisecond separation, the timbre of the ensemble does not change any further, but if presented binaurally, localization effects still occur




Figure 11. Doublet response (echogram) of the New York Philharmonic Hall before reconstruction. This response was collected by D. Gottlob. Note that even though certain distance reflections are clear, there is a great deal of activity throughout the response.

down to a spacing of tens of microseconds. With this in mind, it is clear that if we simulate the discrete echos down to a limit of one millisecond before switching to a recirculating delay reverberator, we can capture much of the effect of a specific concert hall, except possibly for localization effects.

Before we leave the topic of computer simulation of geometric acoustics, we should mention at least two other interesting results of these simulations. First of all, one might criticize the approach for being two-dimensional. However, the formulas do not get terribly more complex when three-dimensional geometry is used, so one could write the simulation program entirely in three dimensions. The winged-edge polyhedra model of Baumgart (see References) seems almost ideally suited to this in that the faces are canonically ordered, making inside/outside determination simple, and that local connectedness is represented simply.

Another interesting result is that non-parallel walls do not necessarily imply uniform echo patterns. Figure 8b, for instance, shows a simulated hall with non-parallel walls that exhibits a highly regular impulse response, regardless of source placement. The phantom sources cluster in circles about the source of radii roughly twice the length of the hall. Likewise, the fact that the best-liked concert halls are exactly rectangular in shape (Vienna Grosser Musikvereinsaal, Boston Symphony Hall, Amsterdam Concertgebouw) leads one to wonder indeed if parallel walls really do have anything to do with the subjective characteristics of the halls. It would seem that the detailed structure and composition of the walls are perceptually more important than the overall shape.

So Now What's Wrong?

We have described several different approaches to simulating concert hall reverberation. One approach is largely

disconnected from physical reality in the use of recirculating delays to simulate the echo patterns of concert halls. The other approach is a computer simulation of the actual reflection pattern of a concert hall based upon its geometry.

The only problem with simulating room reverberation in this way, that is with direct geometrical simulation of the room acoustics, is that it doesn't sound at all like real rooms. This somewhat startling fact means that one or more of our assumptions does not hold water. To follow this up further, we obtained copies of the responses of many natural concert halls which were taken by Shroeder *et al.* (1974) and Gottlob (1975). Careful examination of these impulse responses made it quite clear that a very important consideration that

Table 3. Tap times (in seconds) and gains for a reasonable-sounding seven-tap FIR section for reverberation. Tap number corresponds to the subscript of the coefficient (a_j) and delay (m_j) shown in Figure 12.

TAP	TIME	GAIN
0	0	1.00
1	.0199	1.02
2	.0354	.818
3	.0389	.635
4	.0414	.719
5	.0699	.267
6	.0796	.242


Tap times (in seconds) and gains for a reasonable-sounding 19-tap FIR section for reverberation. These were taken from a highly idealized geometric simulation of the Boston Symphony Hall.

TAP	TIME	GAIN
0	0	1.0
1	.0043	.841
2	.0215	.504
3	.0225	.491
4	.0268	.379
5	.0270	.380
6	.0298	.346
7	.0458	.289
8	.0485	.272
9	.0572	.192
10	.0587	.193
11	.0595	.217
12	.0612	.181
13	.0707	.180
14	.0708	.181
15	.0726	.176
16	.0741	.142
17	.0753	.167
18	.0797	.134

we have been ignoring is the effect of *diffusion*. This comes about when walls are not in fact flat but are irregular. This is, of course, true in all respected concert halls. Boston Symphony Hall has fluted side walls and a box well ceiling, each of which furnishes a plethora of new phantom sources. Figure 11 shows the doublet response of the old New York Philharmonic Hall before reconstruction, obtained by placing a powerful spark gap device on stage and an omnidirectional microphone in a particular seat in the audience. We notice immediately that some distinct echos are visible, but a great confusion soon follows with virtually no silent portions. This confusion is clearly the result of the infinite multiplicity of diffused sources caused by every little irregularity in the hall. The rough surfaces completely wash out all but the first few images, making any use of geometric acoustical modeling, either with phantom sources or ray models, useless for determining the actual sound of a concert hall. These simulations may still be useful for determining reverberation time and other more gross qualities of a given room, but not for the detailed calculation of the impulse response. We can, of course, convolve our music directly with the recorded response (corrected for the spectrum of the spark gap itself) to produce the sound of that concert hall, but even with a 25 KHz sampling rate and a 2.0 second reverberation time,

this comes out to 50,000 multiplications and additions per output sample. This is clearly only in the domain of very patient researchers with great quantities of computer time available. Even with fast Fourier transform techniques (Stockham 1966), we consumed about 6 minutes of computer time (PDP-10) for every second of sound time thus processed. We might then ask if there is any way that we can simulate at least some of these effects without such an obese expenditure of computer time. Indeed, it would appear that a modification of the Schroeder design of Figure 3 can give somewhat improved results at only a moderate increase in computation time.

Figure 12 shows our new formulation. As before, we have N delays which simulate the first N reflections plus the direct signal. This presumably takes care of the first 40 to 80 milliseconds of the decay. This entire signal is now forwarded to the reverberator, $R(z)$, which is then added into the early decay with a gain of g . What this provides is that all the density of the first N reflections is then repeated many many times by the recirculating delays of the reverberator. The delays D_1 and D_2 are set such that the first echo from the reverberator coincides with the end of the last echo from the early response. This means that either D_1 or D_2 will be zero, depending on whether the total delay



of the early echo is longer or shorter than the shortest delay in the reverberator.

Adding this separate discrete early echo makes a substantial difference in the sound, and in fact, if the actual room response is used for the early delays, the sound does begin to approach the sound of a real concert hall. Since using the first 80 milliseconds of the actual room impulse response is still well beyond the bounds of sanity, except in certain restricted circumstances for research purposes, we may then ask how simple can this early response be and still give a rich, life-like sound.

Various ways were tried to synthesize a suitable early response, from using the results of the geometric simulation of a room, real or fictitious, to choosing random numbers to determine the positions and amplitudes. Many problems became clear while attempting to do this. It would seem that one cannot just compute a bunch of impulses at random and expect the result to sound good, in that there is always the possibility that a very strong comb filter effect will be formed. Even if the positions are chosen to be prime numbers of samples from the origin, the differences among them in time may well be highly composite. This is easy to see when one realizes that the sum or difference of two prime numbers is always divisible by two. As a result, we cannot give at this time any general method for finding different early responses that will sound good. We have, essentially by trial and error, determined at least two that sound reasonable without too much coloration of the sound. These are given in Table 3. There is a seven-tap section and a nineteen-tap section. The first tap is the direct signal and can be changed to provide more or less reverberation if desired.

Thus, we are now in a position to give a systematic technique for designing a reasonably good-sounding reverberator, albeit not one that corresponds to any particular concert hall in existence today. The formulation of Figure 12 with the nineteen-tap model is preferred, although the seven-tap section can be used, or even a no-tap system if computer time is at a premium. The recirculating part, as mentioned before, should consist of six combs in parallel, each with a first-order lowpass filter in the loop. The outputs of the six combs should be summed and forwarded to a single all-pass network with a gain of 0.7 and a delay of about six milliseconds. The gains of the low-pass filters in the feedback loops of the combs should be set to the values from Figure 6 for the delay used. The delays should be linearly distributed over a ratio of 1:1.5 with the shortest delay being about 50 milliseconds. Again, if computer memory is at a premium, this can be reduced to as low as ten milliseconds without gross degradation. The delay lengths in samples should be set to the closest prime numbers to prevent exactly overlapping pulses, although with higher values of g_1 in the feedback filters, the sensitivity to the exact length of the delay becomes negligible. The gains g_2 of each comb should be set to $g(1-g_1)$ where g now determines the overall reverberation time. Figure 13 shows the relation between total reverberation time as a function of g , computed at a sampling rate of 25 KHz.

We originally computed the reverberation time both as twice the time to fall from 5 to 35 dB, and as the time to fall from 0 to 60 dB, but these numbers were quite similar (within 10%) at most values of g except very small or very large ones. The reverberation time of this system can be estimated without use of Figure 13 as $.366/(1-g)$ to an error

of about 10%. Likewise, the gain can be computed for a given value of reverberation time as $1 - .366/T$ where T is the desired reverberation time in seconds. Note that this will change with sampling rate, but the deviation should not be too drastic. Likewise, this time will scale directly with the delay lengths, as long as the 1:1.5 ratio is maintained between the shortest comb delay and the longest.

We should note here one somewhat striking discovery we made in simulating natural sounding impulse responses. While digitizing the impulse responses from concert halls around the world, we kept noticing that the responses in the finest concert halls sounded remarkably similar to white noise with an exponential amplitude envelope. To test this observation, we generated synthetic impulse responses by shaping unit-variance Gaussian pseudo-random sequences with an exponential of the desired length. The direct sound was added by including an impulse at the beginning. The balance between the impulse height and the strength of the decay was determined empirically to be about 12:1 for a natural sound. We then convolved this synthetic impulse response with a variety of unreverberated musical sources and compared the results with the sounds of actual concert halls. The results were astonishing. Although the synthetic impulse response did not produce a sound that could be identified with a specific concert hall, the sound was clearly a very natural-sounding response. With this method of creating synthetic impulse responses, one can "tailor" the characteristics to one's needs. By selectively filtering the impulse response before convolution, we can get any desired rolloff rates at various frequencies. It was the most natural-sounding of all the techniques we have tried, outside of using the measured impulse response of a concert hall. Stereo can then be produced by simply generating two synthetic impulse response sequences and calculating the left and right ears separately.




Figure 13. Reverberation time of the six-comb, one-allpass reverberator, assuming that the low-pass filter gain, g_1 , is set according to the guides of Figure 6 and the delays are set according to Table 2. In this case, the gain g will determine the reverberation time. This curve can be approximated by the function $.366/(1-g)$. Conversely, g can be estimated as $1 - .366/T$ where T is the reverberation time desired (in seconds).

This gives perfectly uncorrelated binaural reverberation. Frequency dependence in the decay can easily be introduced by filtering the synthetic impulse response before doing the convolution. Again, the only problem with this is the enormous amount of computation involved for direct convolution. The fast Fourier transform can help ease the computational burden, but real-time operation is still more than a factor of ten away for even the fastest commercially available signal processors. One observation that we can offer is that the impulse response decays fairly rapidly such that less than 30% of the points need more than twelve bits of precision. Another 30% need between eight and eleven bits, and around 10% can be specified in four bits or less. This all assumes a sixteen-bit sample width.

Conclusions and Suggestions for Further Work

In the course of this work, it became quite clear that all the geometric simulations of concert hall acoustics that have been done to date result in a simulated room reverberation that does not sound at all like real rooms. This is probably due to the effect of diffusion of echos due to irregularities in the reflecting surfaces, and to the fact that the spectrum of the echo is modified by the reflection in a manner that depends strongly on the angle of incidence.

Despite these facts, we have been able to achieve a good-sounding, smooth artificial reverberation system that, regardless of the fact that it does not sound like actual concert halls, does sound good enough to be used everywhere that artificial reverberation is desired, and does eliminate some of the problems that plagued earlier reverberation schemes, such as "fluttery" response to short, impulsive sounds and metallic-sounding decay. We also found a much larger number of non-useful unit generators than useful new unit generators.

There are several topics that were revealed in the course of this study that would make excellent material for further study, and excellent doctoral dissertation topics. These might include the following:

- 1) Improve the computer modeling of concert halls by geometry. This would require a leap of insight to figure out a reasonable way of dealing with the effects of diffusion, or even a way of describing the diffusive properties of reflecting surfaces.
- 2) Figure out what gives recirculating delay reverberators their characteristic "sound." This would be an acoustic/psychoacoustic study in timbre. One would have to vary the physical characteristics of the impulse response and attempt to correlate this with the subjective perception, possibly involving the use of multi-dimensional scaling or factor analysis to help decompose the perceptual dimensions.
- 3) Make up some new unit reverberators and see what they sound like. Try putting more complex filters in the loop and see if that has any substantial effect (it might!).
- 4) Figure out some clever way to do all this with less computer time, possibly using the whasit transform, or the whoosit decomposition.

Acknowledgements

As usual, it would be out of the question to attempt to thank individually all the people and institutions that went into this research. My thanks, as always, go to my dear colleagues at Stanford: J. Chowning, J. Grey, G. Loy, L. Rush, and P. Wood. Thanks to D. Gottlob who so patiently sat through two harrowing days of digitizing all these impulse responses with us at IRCAM. Thanks also to Jean Derechapt, my co-conspirator and programming aid in this work, who got my imagination going on this subject. Finally, of course, to the Institut de Recherche et Coordination Acoustique/Musique (commonly known as IRCAM) for providing one of the few places in the world where this kind of in-dept research can be done, and to my boss, Gerald Bennett, for letting me heed my notions in plowing through the maze of half-baked ideas.

References

- Baumgart, B.G. "Winged Edge Polyhedron Representation." *Stanford Artificial Intelligence Laboratory Memo* Number AIM-179, Stanford University, October, 1972.
- Baumgart, B. G. "GEOMED - A Geometric Editor." *Stanford Artificial Intelligence Laboratory Memo* Number AIM-232, Stanford University, May, 1974.
- Baumgart, B.G. "Geometric Modeling for Computer Vision." *Stanford Artificial Intelligence Laboratory Memo* Number AIM-249, Stanford University, October, 1974.
- Baumgart, B. G. "A Polyhedron Representation for Computer Vision." *Proceedings of the National Computer Conference*, 1975, pp. 589-596.
- Beranek, L. L. *Acoustics*. New York: McGraw-Hill, 1954.
- Beranek, L. L. *Music, Acoustics, and Architecture*. New York: John Wiley and Sons, 1962.
- Date, H., and Y. Tozuka. "An Artificial Reverberator whose Amplitude and Reverberation Time-Frequency Characteristics can be Controlled Independently." *Acoustics*, 17:42-47, 1966.
- Gottlob, D. *Vergleich objektiver akustischer Parameter mit Ergebnissen subjektiver Untersuchungen an Konzertsälen*. Ph.D. Dissertation, Georg-August-Universität, Göttingen, Germany, 1975.
- Green, D. M. "Temporal Auditory Acuity." *Psychological Review*, 78(6):540-551, 1971.
- Harris, C. M. "Absorption of Sound in Air in the Audio-Frequency Range." *Journ. Acoust. Soc. Amer.*, 35:11, 1963.
- Harris, C.M., and W. Tempest. "Absorption of Sound in Air below 1000 cps." *Journ. Acoust. Soc. Amer.*, 36:2390, 1964.
- Kuttruff, H. *Room Acoustics*. London: Applied Science Publishers, 1973.
- Marquardt, D. W. "An Algorithm for Least-Squares Estimation of Non-Linear Parameters." *Siam Journal*, 11(2): 431-441, 1963.
- Oppenheim, A.V., and C.J. Weinstein. "Effects of Finite Register Length in Digital Filtering and the Fast Fourier Transform." *Proceedings of the IEEE*, 60(8):957-976, 1972.

- Oppenheim, A.V., and R. W. Schafer. *Digital Signal Processing*. Englewood Cliffs, New Jersey: Prentice-Hall, 1975.
- Sabine, W. C. "Reverberation," in *Acoustics: Historical and Philosophical Development*, ed. by R. B. Lindsay. Stroudsburg, Pennsylvania: Dowden, Hutchinson, and Ross, 1972.
- Schroeder, M.R. "Improved Quasi-Stereophony and Colorless Artificial Reverberation." *Journ. Audio Eng. Soc.*, 33:1061, 1961.
- Schroeder, M.R. "Natural-Sounding Artificial Reverberation." *Journ. Audio Eng. Soc.*, 10:219-223, 1962.
- Schroeder, M.R. "Digital Simulation of Sound Transmission in Reverberant Spaces." (Part 1). *Journ. Acoust. Soc. Amer.*, 47(2):424-431, 1970.
- Schroeder, M.R., D. Gottlob and K.F. Siebrasse. "Comparative Study of European Concert Halls." *Journ. Acoust. Soc. Amer.*, 56:1195-1201, 1974.
- Stockham, T.G., Jr. "High-Speed Convolution and Correlation." *AFIPS Conference Proceedings*, 1966 Spring Joint Computer Conference, 28:229, 1966.

continued from page 4

CONFERENCE ON ARTIFICIAL INTELLIGENCE

The Sixth International Joint Conference on Artificial Intelligence will be held in Tokyo, from August 20 to 24, 1979. For more information, contact Prof. Bruce Buchanan, Computer Science Department, Stanford University, Stanford, California 94305, phone (425) 497-4878.

Computer Music Journal

is available in
microform



Please send me additional information.

University Microfilms International

300 North Zeeb Road
Dept. P.R.
Ann Arbor, MI 48106
U.S.A.

18 Bedford Row
Dept. P.R.
London, WC1R 4EJ
England

Name _____
Institution _____
Street _____
City _____
State _____ Zip _____

"COMPUTERS AND ENTERTAINMENT"

The 1980 National Computer Conference (to be held May 19-22, 1980, at Anaheim, California) will focus on research and application of computer technology in the following areas:

- Movies and Television
- The Performing Arts (Music, Dance, Stage)
- Home Entertainment
- Amusements (Parks, Zoos, Casinos, Special Theatres)
- Sports
- Art
- Literature

You are invited to participate in this special program by:

- Writing and presenting a paper;
- Proposing a session;
- Volunteering to be a panelist.

All papers are referred and sessions are subject to the approval of the Program Committee. The NCC does not grant honorariums, nor does it pay for travel, food, or lodging expenditures. The deadline for all submissions is November 1, 1979. However, notification of your intentions as early as possible would be appreciated. For more information, contact: Suzanne Landa, Information Sciences Department, The Rand Corporation, 1700 Main Street, Santa Monica, California 90406. (213) 393-0411, Ext. 7171.

LA JOLLA CONFERENCE ON COGNITIVE SCIENCE

An international conference on Cognitive Science will be held from August 13 to 16, 1979 on the campus of the University of California, San Diego. An attempt has been made to select speakers and symposia over a broad range of Cognitive Science. Speakers will include: N. Geschwind, P. Johnson-Laird, G. Lakoff, M. Minsky, A. Newell, D. Norman, R. Schank, J. Searle, H. Simon, and T. Winograd.

For additional information, contact Nomi Feldman, 2321 Morena Blvd., San Diego, California 92110, phone (714) 275-0650.

continued from page 11

References

1. Aho, A. V., et al. (1977) "Code Generation for Expressions with Common Subexpressions," *Journal of the Association for Computing Machinery*, Vol. 24, No. 1, Jan. This would be a good starting point for exploring the contributions of ordinary code generation theory to a partially parallel model.
2. Aho, A. V., and Ullman, J. D. (1978) *Principles of Compiler Design*, Addison-Wesley. This is a good text on compiling, including code generation and optimization techniques. It also includes a discussion of data structures appropriate to some of the techniques used.
3. Baker, Kenneth R. (1974) *Introduction to Sequencing and Scheduling*, Wiley and Sons, New York.
4. Bruno, J. L., and Sethi, R. (1976) "Code Generation for a One-register Machine," *Journal of the Association for Computing Machinery*, Vol. 23, No. 3, July.
5. Thesen, Arne (1976) "Heuristic Scheduling of Activities Under Resource and Precedence Restrictions," *Management Science*, Vol. 23, No. 4, Dec. This article discusses an algorithm for a scheduling problem which can be adapted to look like the one discussed above, and gives pointers to other literature.