

(WDR) built a motorized sound rotation system for concert performance of Stockhausen's works (Morawska-Büngler 1988).

Simulation of Rotating Loudspeakers

The effects of rotation are manifold, involving Doppler shift vibrato, time-varying filtering, phase shifts, distortions caused by air turbulence, and echo reflections from adjacent surfaces—not to mention the transfer characteristics of the amplifiers and loudspeakers used. The Leslie Tone Cabinet, for example, employed vacuum tube electronics with “overdrive” distortion if desired. These complicated and interacting acoustical and electronic effects are difficult to simulate convincingly using digital signal processing. Nonetheless, a number of synthesizers and effects units offer programs that simulate rotating loudspeakers. Such programs should improve as more sophisticated algorithms are developed.

Reverberation

Reverberation is a naturally occurring acoustical effect. We hear it in large churches, concert halls, and other spaces with high ceilings and reflective surfaces. Sounds emitted in these spaces are reinforced by thousands of closely spaced echoes bouncing off the ceiling, walls, and floors. Many of these echoes arrive at our ears after reflecting off several surfaces, so we hear them after the original sound has reached our ears. The ear distinguishes between the *direct* (original) sound and the *reflected* sound because the reflected sound is usually lower in amplitude, slightly delayed, and lowpass filtered due to absorption of high frequencies by the air and reflecting surfaces (figure 11.15). The myriad echoes fuse in our ear into a lingering acoustical “halo” following the original sound.

A microphone recording of an instrument in a concert hall is surrounded by an envelope of reverberation from the hall. This is particularly the case when the microphone has an omnidirectional pattern. For recordings made in small studio spaces, it is often desirable to add reverberation, since without it a voice or ensemble sounds “dry,” lacks “space” or “depth.”

Certain synthesized sounds have little or no intrinsic spaciousness. These acoustically “dead” signals can be enhanced by spatial panning, echoes, and reverberation processing.

But space is not merely a cosmetic appliqué for sounds. Spatial depth can be used to isolate foreground and background elements in a compositional architecture. Further, reverberation is not a monolithic effect; there are

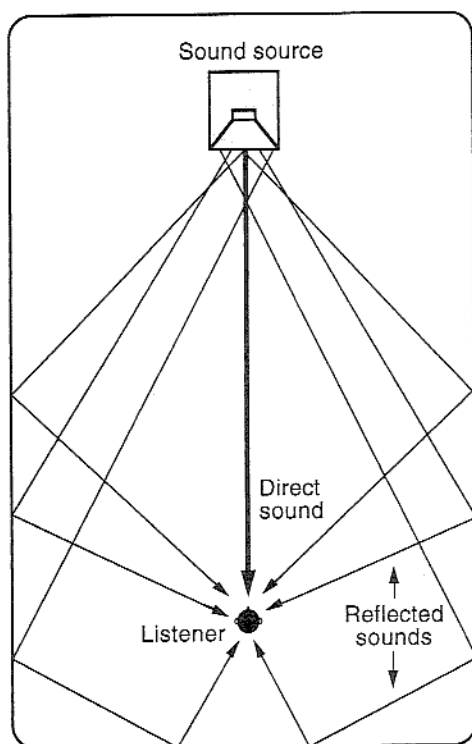


Figure 11.15 Reverberation is caused by reflections of sound off surfaces in a space. The dark line is the path of direct sound; all other lines represent sonic reflections that arrive later than the original due to their longer paths.

many colors and qualities of reverberation—as many as there are natural spaces and synthetic reverberators. No single type of reverberation (natural or synthetic) is ideal for all music. Most electronic reverberation units simulate several types of reverberation. Some attempt (often crudely) to simulate known concert halls, while others create bizarre spatial images that would be impossible to duplicate in a real hall.

Properties of Reverberation

Glorious-sounding salons and concert halls have been constructed since antiquity, but their basic acoustical properties were not well understood from a scientific standpoint until the late nineteenth century. The pioneering work on the analysis of reverberant spaces was carried out by Wallace Sabine (1868–1919), who advised in the construction (starting from an existing structure) of Boston's acclaimed Symphony Hall in 1900. Symphony Hall was the first performance space designed according to rigorous,

scientifically derived principles of acoustics. Sabine observed that a room's reverberation is dependent on its volume, geometry, and the reflectivity of its surfaces (Sabine 1922). It is no surprise that large rooms with reflective surfaces have long reverberation times, and small rooms with absorptive surfaces have short reverberation times. Smooth, hard surfaces like glass, chrome, and marble tend to reflect all frequencies well, while absorptive surfaces like heavy curtains, foam, and thick carpeting tend to absorb high frequencies.

The geometry of the room surfaces determines the angle of sound reflections. Walls that are not parallel scatter the wavefronts in complicated dispersion patterns, and small irregularities such as plaster trimmings, indentations, columns, and statues tend to diffuse the reflections, creating a richer, denser reverberation effect.

Sabine also observed that humidity affects the reverberation time in large halls, since up to a point, humid air tends to absorb high frequencies.

Impulse Response of a Room

One way to measure the reverberation of a room is to trigger a very short burst (an *impulse*) and plot the room's response over time. This plot, when corrected for the spectrum of the burst, shows the *impulse response* of the room. As mentioned in chapter 10, circuits also exhibit an impulse response, making the impulse response measurement an often-used tool both in circuit design and concert hall design. Natural reverberation typically has an

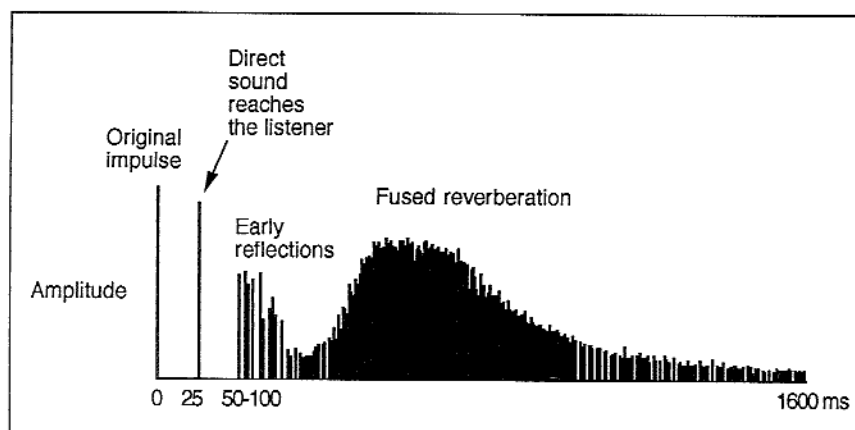


Figure 11.16 Impulse response envelope of a reverberant hall. The components of reverberation are shown as the *predelay* (shown as a 25-ms delay before the direct sound reaches the listener), the early reflections, and the fused reverberation.

impulse response envelope similar to that shown in figure 11.16. The build-up of reverberation follows a quasi-exponential curve that reaches a peak within a half-second and decays more or less slowly.

In general, an irregular time interval between peaks is desirable in a concert hall. Regularly spaced peaks indicate “ringing”—resonant frequencies in the hall—which can be annoying.

Reverberation Time

Another important measurement of reverberation is *reverberation time* or RT60. The term RT60 refers to the time it takes the reverberation to decay 60 dB from its peak amplitude (1/1000 of its peak energy). Typical RT60 times for concert halls are from 1.5 to 3 seconds. The RT60 point of the plot in figure 11.17 is 2.5 seconds.

Artificial Reverberation: Background

The earliest attempts at artificial reverberation of recordings transmitted the sound through an *acoustic echo chamber*, then mixed the reverberated signal with the original. Some large recording studios still allocate a separate room as an echo chamber. They place a loudspeaker at one end of a *reflective* room and put a high-quality microphone at the other end. The sound to be reverberated is played over the loudspeaker and picked up by the microphone (figure 11.18). An echo chamber offers a unique acoustical *ambience* created by a specific room, loudspeaker, and microphone. When all these conditions are sympathetic, the quality of reverberation may be excellent. A drawback to the echo chamber approach (besides the practicalities of

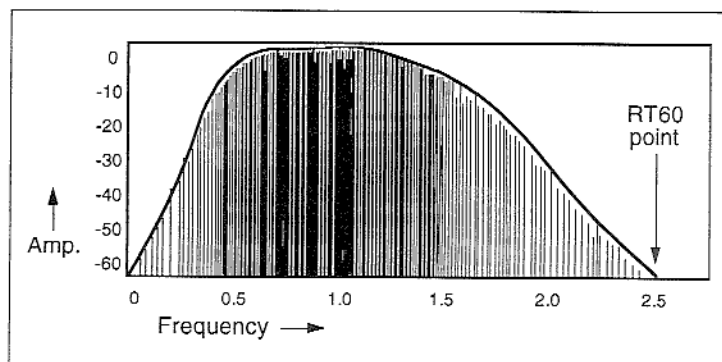


Figure 11.17 Reverberation time is measured as the point at which the reverberation decays to -60 dB of its peak level.

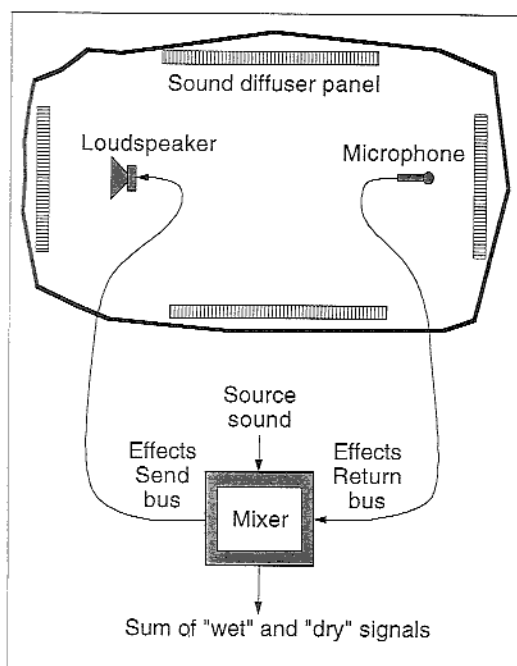


Figure 11.18 To create an acoustic ambience effect, sound can be fed into an echo chamber via a loudspeaker. The reflected, indirect sound is picked up by a microphone at the other end of the room. Ideally, the room is irregularly shaped. To maximize and randomize the reflections, the room should be fitted with *sound diffuser panels*. Sound diffuser panels contain many recesses spaced at different distances. As sound waves strike them they are reflected back at different delay times depending on which recess they hit. This diffusion effect tends to eliminate *standing waves* (resonant frequencies in the room) caused by parallel walls.

constructing such a space) is that the reverberation cannot be varied tremendously.

The more usual way of adding reverberation is with a *reverberation unit* or *reverberator*. Before digital reverberators were introduced in the mid-1970s, reverberators were electromechanical contraptions containing two transducers (input and output) and a reverberating medium like a long metal spring or a metal plate. The sound to be reverberated was transmitted from the input transducer to the medium. The medium transmitted the sound to the output transducer mixed with myriad echoes caused by vibrations/reflections of the signal within the medium. The result was amplified and mixed with the original signal to create a rather "colored" artificial reverberation effect. The best plate reverberators produced relatively clean and diffuse reverberation, but they were limited to an RT60 of only a few seconds and a fixed reverberation pattern.

Digital Reverberation Algorithms

Digital reverberators use time delays, filters, and mixing to achieve the illusion of sound scattering within a room. From a signal-processing standpoint, a reverberator is a filter with an impulse response that resembles the impulse response of a room. Manfred Schroeder of the Bell Telephone Laboratories (1961, 1962, 1970) was the first to implement an artificial reverberation algorithm on a digital computer. His reverberation programs soaked up hours of computation time on the behemoth mainframe computers of the epoch. Modern reverberation units are compact and run in real time. Control knobs and buttons on their front panels let musicians dial up a variety of effects. Most reverberators can be controlled via MIDI (see chapter 21).

Parts of Reverberation

The effect of reverberation can be broken into three parts, shown earlier in figure 11.16.

- *Direct (unreflected) sound* travels in a straight path and is the first sound to arrive at the listener's ears
- *Discrete early reflections* hit the listener just after the direct sound
- *Fused reverberation* contains thousands of closely spaced echoes but takes some time to build up and then fade away

Commercial reverberation units usually provide controls that let one manipulate these parts more or less independently. On these units, the balance between the reverberated and direct sound is sometimes called the *wet/dry* ratio (the reverberated sound is said to be “wet”), and the delay just before the early reflections is called the *predelay*.

Effective simulation of natural reverberation requires high *echo density*. Some early digital reverberators produced as few as 30 echoes per second, while in actual concert halls, an echo density of more than 1000 echoes per second is typical. Many reverberators today provide a control that lets users adjust the echo density to suit the desired effect, from discrete echoes to a dense, fused reverberation pattern.

The discrete early reflections of a concert hall can be simulated by means of a *tapped delay line*. This is simply a delay unit that can be “tapped” at several points to put out several versions of the input signal, each delayed by a different amount. (See chapter 10 for an explanation of tapped delay lines.)

The lush sound of fused reverberation requires a greater echo density than a tapped delay line can efficiently provide. Many different algorithms for fused reverberation exist, but they all usually involve a variation on M. R. Schroeder's original algorithms, described next.

Unit Reverberators

Schroeder called the building blocks *unit reverberators*, of which there are two forms: *recursive comb filters* and *allpass filters*, both of which were introduced in chapter 10.

Recursive Comb Filters

As explained in chapter 10, a recursive or *infinite impulse response* (IIR) comb filter contains a feedback loop in which an input signal is delayed by D samples and multiplied by an amplitude or gain factor g , and then routed back to be added to the latest input signal (figure 11.19a).

When the delay D is small (i.e., less than about 10 ms) the comb filter's effect is primarily a spectral one. That is, it creates peaks and dips in the frequency response of the input signal. When D is larger than about 10 ms, it creates a series of decaying echoes, as shown in figure 11.19b. The echoes decay exponentially, so for the maximum number of echoes (the longest decay time), g is set to nearly 1.0. The time it takes for the output of the comb filter to decay by 60 dB is specified by the following formula (Moore 1990):

$$\text{decay_time} = (60 / -\text{loopGain}) \times \text{loopDelay}$$

where loopGain is the gain g expressed in decibels $= 20 \times \log_{10}(g)$, and loopDelay is the delay D expressed in seconds $= D/R$, where R is the sampling rate. Thus if $g = 0.7$, then $\text{loopGain} = -3$ dB.

Allpass Filters

Allpass filters transmit all frequencies of steady-state signals equally well (see chapter 10). But they "color" sharp transient signals by introducing frequency-dependent delays. When the delay time is long enough (between 5 and 100 ms), the allpass filter shown in figure 11.20a has an impulse response as shown in figure 11.20b: a series of exponentially decaying echo pulses, like a comb filter with a long delay. The uniform spacing between the pulses suggests that when a short, transient sound is applied, the filter rings with a period equal to the delay time of the filter. This explains why allpass

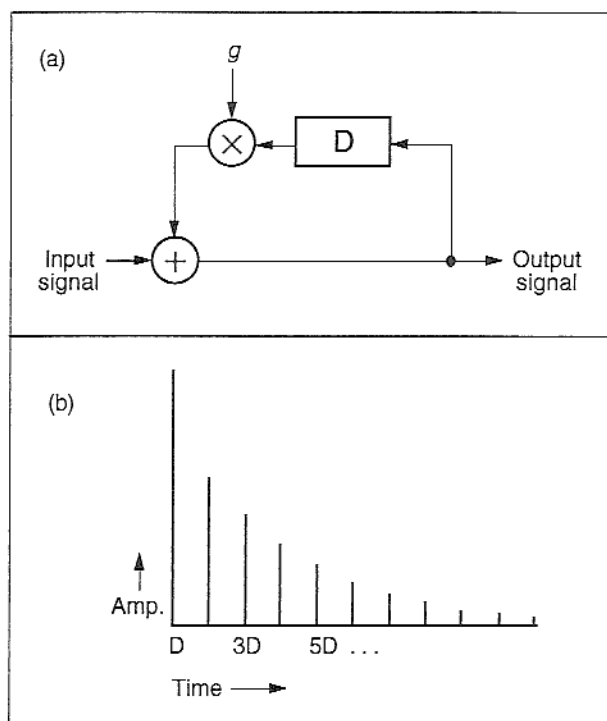


Figure 11.19 A recursive comb filter for reverberation. (a) Circuit of comb filter with coefficients D (number of samples to delay) and g (amount of feedback). (b) Impulse response, as a series of echoes.

filters are not “colorless” when they treat sounds with sharp attack and decay transients.

Reverberation Patches

We have established that both recursive comb and allpass filters can generate a series of decaying echoes. For lush reverberation, it is necessary to interconnect a number of unit reverberators to create sufficient echo density so that the echoes fuse. When unit reverberators are connected in parallel, their echoes add together. When they are connected in series, each echo generated by one unit triggers a series of echoes in the next unit, creating a much greater echo density. The number of echoes produced in series is the product of the number of echoes produced by each unit.

In Schroeder’s designs, comb filters are interconnected in parallel to minimize spectral anomalies. For example, a frequency that passes through one comb filter might be attenuated by another. Allpass filters are usually connected in series. Because of the phase distortion they introduce, connecting

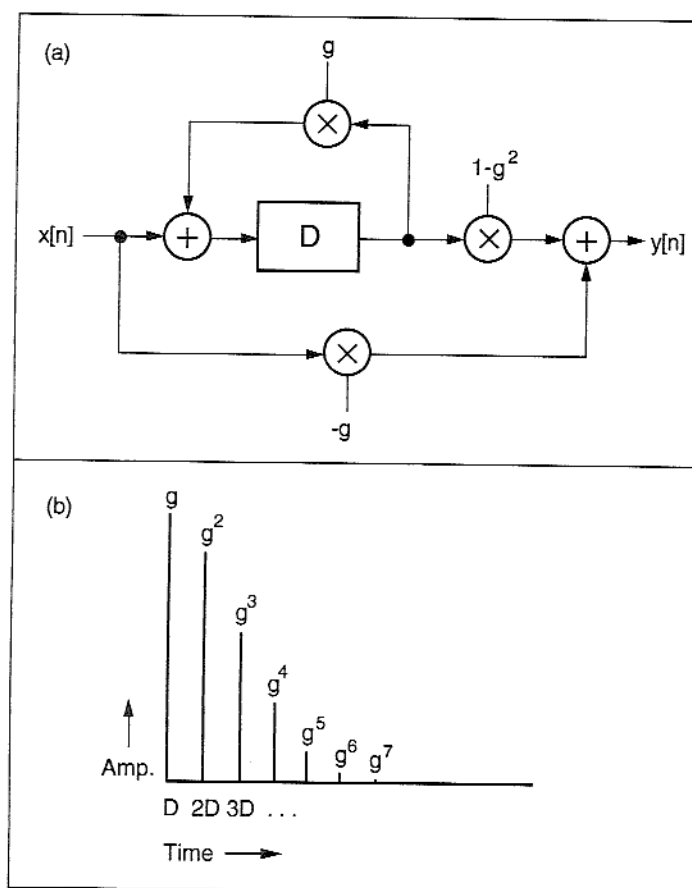


Figure 11.20 A first-order allpass network. (a) By adding $-g$ times the input into the output of the delay, a comb filter is changed into an allpass filter. (b) The impulse response of an allpass filter is an exponentially decaying series of echo pulses. This makes the impulse filter useful as a building block of reverberators.

allpass filters in parallel can result in a nonuniform amplitude response due to phase cancellation effects.

Figure 11.21 shows two reverberators proposed by Schroeder. In figure 11.21a the parallel comb filters initiate a train of echoes that are summed and fed to two allpass filters in series. In figure 11.21b five allpass filters cause the echo density to be multiplied by each unit. If each allpass generates just four audible echoes, the end result is 1024 echoes at the output of allpass number 5.

The characteristic sound of a digital reverberation system of this type is dependent on the choice of the delay times D (these determine the spacing of the echoes) and amplitude factors g (these determine the decay or rever-

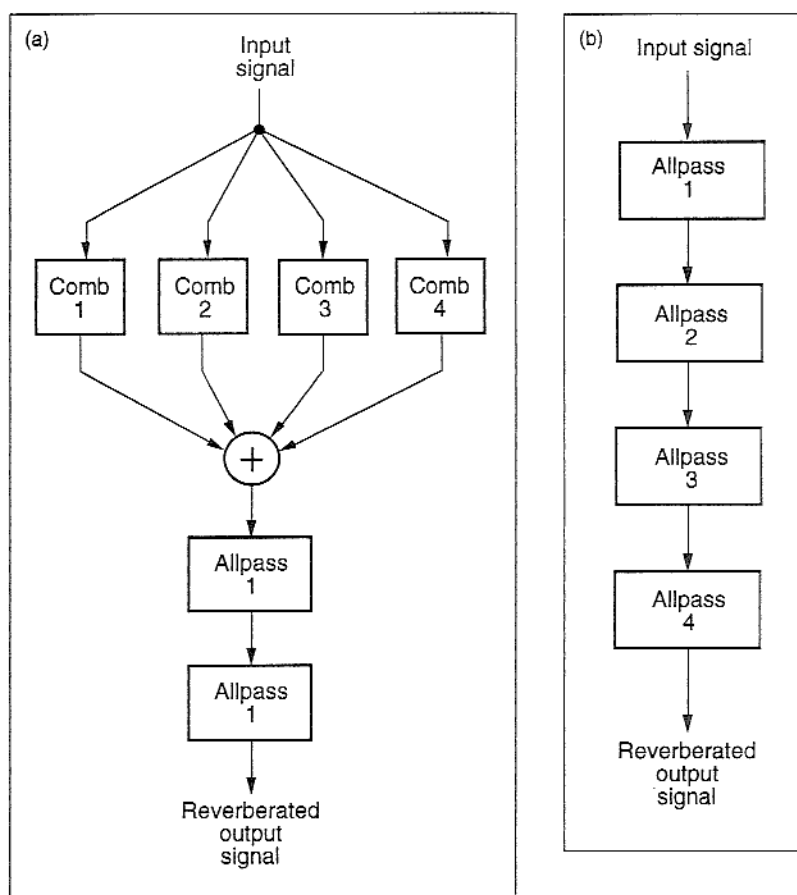


Figure 11.21 Schroeder's original reverberator designs. (a) Parallel comb filters fed into two allpass filter stages. (b) Five allpass filter stages in series.

beration time) for each of the unit reverberators inside it. The delay time is also called the *loop time*.

For natural-sounding reverberation it is important to choose delay times that are relatively prime to one another (i.e., that have no common divisor) (Moorer 1977, 1979c). Why is this? Consider two comb filters, where the delay time of the first is 10 ms and that of the second is 12.5 ms. The length of their delay lines are 800 samples and 1000 samples, respectively, at a sampling rate of 40 KHz. Because the lengths of both delay lines are divisible by 200, a reverberator built from these two units does not have a smooth decay. At multiples of 200 ms, the echoes coincide to increase the amplitude at that point, causing a sensation of discrete echoes or regular "bumps" in the decay. When the delay times are adjusted to 10.025 and 24.925 ms, the length of their delay lines are 799 and 997, respectively. Now

the first coincidence of echoes does not occur until $(799 \times 997)/40 \text{ KHz} = 19.91$ seconds. (See Moorer 1979c for a discussion of how to tune these parameters.)

As might be expected, shorter delay times correlate with the sound of smaller spaces. For a large concert hall, the reverberator in figure 11.21a uses comb filter delay times around 50 ms with a ratio of longest:shortest delay of 1.7:1. For a small tiled room effect the comb filter delay times can be set in the range of 10 ms. The allpass filters have relatively short loop times of 5 msec or less. The reverberation time of the allpass filters must be short (less than 100 msec) because their purpose is to increase the density of the overall reverberation, not its duration.

Simulation of Early Reflections

Schroeder's reverberation algorithms can be characterized as *tapped recirculating delay* (TRD) models. As explained earlier, the reverberator is usually partitioned into comb and allpass sections, which generate sufficient echo density to create a reasonable simulation of *global reverberation*. The TRD model is efficient, but it simulates only generic global reverberation, and not the detailed acoustic properties of an actual performance space.

In 1970 Schroeder extended his original reverberator algorithms to incorporate a *multitap delay line* to simulate the early reflections that are heard in a hall before the onset of the fused reverberant sound. (See chapter 10 for more on multitap delay lines.) This design, which has been adopted in most commercial reverberators, is shown in figure 11.22. Thus to simulate a particular concert hall, a straightforward way to improve the basic TRD model is to graft the measured early reflection response of the hall onto the generic global reverberator (Moorer 1979c). A further extension is to lowpass filter the global reverberation according to the measured sound absorption characteristics of the hall.

Another important consideration in reverberation design is that the sounds presented in each ear should be *mutually incoherent*. That is, the reverberation algorithm should be slightly different (*decorrelated*) for each channel of processing.

Fictional Reverberation Effects

The goals of the electronic music composer extend beyond the simulation of natural reverberant spaces. A reverberator can conjure up many unusual "fictional" spatial effects that are not meant to be realistic. A common example is "gated" reverberation that explodes quickly in echo density,

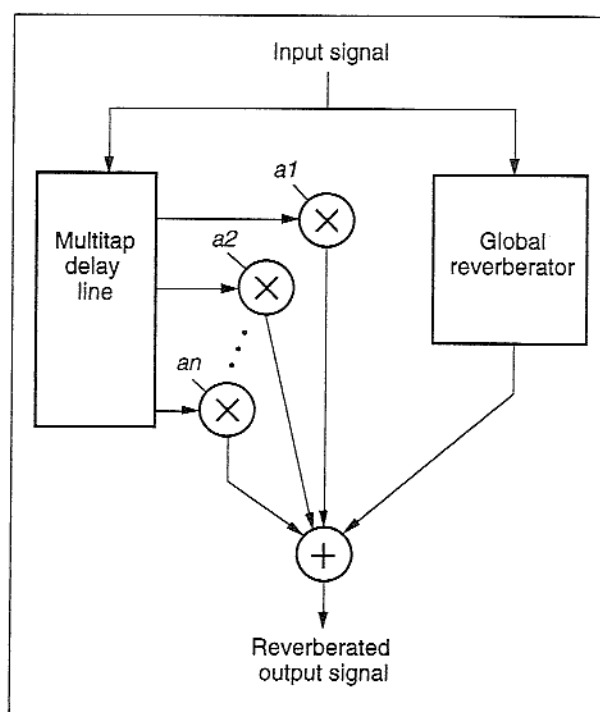


Figure 11.22 In Schroeder's later designs, a multitap delay line simulated the *early* reflections of sound in a concert hall.

Table 11.2 Typical parameters of reverberators

Parameter	Description
Type of reverberation	Choice between "Hall," "Chamber," "Plate," or "Gated"
Size	Sets the delay times within the unit reverberators
Predelay	Controls the onset time of the effect
Input delay	Causes the effect to precede the cause (the wet sound precedes the dry sound)
Reverberation time	Sets the decay time
Diffusion	Determines the echo density
Mix	Ratio of input sound to reverberated sound at the output of the device
Highpass filter	Reverberates only the upper octaves of the sound, creating a "sizzling" reverberation effect
Lowpass filter	Reverberates only the lower octaves of the sound, creating a "muffled" reverberation effect

then cuts off suddenly. Gated reverberation was used on snare drums in the 1980s and quickly became a pop music cliché. Other effects include a “sizzling” reverberation, obtained by applying a highpass filter to the reverberated sound, and its opposite, a muffled reverberation obtained by applying a steep lowpass filter. By manipulating the parameters of a reverberator, one can create weird combinations such as tiny rooms with long reverberation times. Table 11.2 lists the parameters provided on many commercial reverberators.

The section on reverberation by convolution, later in this chapter, presents another type of nonrealistic reverberation using the asynchronous granular synthesis technique covered in chapter 5.

Modeling Sound Spaces

The study of reverberation is ongoing. The algorithms described in the earlier section on reverberation are a starting point for the designs discussed here. This section explains several approaches to more realistic reverberation that have been developed in recent years. These include extensions to the basic Schroeder algorithms, geometric models, reverberation via convolution, waveguide reverberation, and multiple-stream reverberation.

Several of these techniques represent a *physical modeling* approach to reverberation. (See chapter 7 for an introduction to the theory of physical modeling in the context of sound synthesis.) These mathematically intensive methods model the diffusion of acoustical waves in actual spaces. Besides creating more realistic models, they offer the possibility of simulating imaginary spaces. In this category we include rooms whose characteristics and geometry can change over time—such as an elastic concert hall that “expands” and “shrinks” over the course of a phrase—or impossible spaces such as a closet with a long reverberation time. Thus the goal of these techniques is not always realistic reverberation, but rather a dramatic spatial transform.

Extensions to Schroeder Reverberation Algorithms

In the standard Schroeder reverberation algorithms, the allpass filters generate a series of echoes with an exponential decay. An extension to the Schroeder model is to substitute an *oscillatory allpass* filter for the regular allpass filter in the Schroeder design. In this case, the impulse response of the allpass filter is a pulse train with an amplitude of a damped sinusoid