

3 REVERBERATION ALGORITHMS

William G. Gardner

MIT Media Laboratory, 20 Ames Street, Cambridge, MA 02139

billg@media.mit.edu

Abstract: This chapter discusses reverberation algorithms, with emphasis on algorithms that can be implemented for realtime performance. The chapter begins with a concise framework describing the physics and perception of reverberation. This includes a discussion of geometrical, modal, and statistical models for reverberation, the perceptual effects of reverberation, and subjective and objective measures of reverberation. Algorithms for simulating early reverberation are discussed first, followed by a discussion of algorithms that simulate late, diffuse reverberation. This latter material is presented in chronological order, starting with reverberators based on comb and allpass filters, then discussing allpass feedback loops, and proceeding to recent designs based on inserting absorptive losses into a lossless prototype implemented using feedback delay networks or digital waveguide networks.

3.1 INTRODUCTION

Our lives are for the most part spent in reverberant environments. Whether we are enjoying a musical performance in a concert hall, speaking to colleagues in the office, walking outdoors on a city street, or even in the woods, the sounds we hear are invariably accompanied by delayed reflections from many different directions. Rather than causing confusion, these reflections often go unnoticed, because our auditory system is well equipped to deal with them. If the reflections occur soon after the initial sound, the result is not perceived as separate sound events. Instead, the reflections modify the perception of the sound, changing the loudness, timbre, and most importantly, the spatial characteristics of the sound. Late reflections, common in very reverberant

environments such as concert halls and cathedrals, often form a background ambience which is quite distinct from the foreground sound.

Interestingly, the presence of reverberation is clearly preferred for most sounds, particularly music. Music without reverberation sounds dry and lifeless. On the other hand, too much reverberation, or the wrong kind of reverberation, can cause a fine musical performance to sound muddy and unintelligible. Between these extremes is a beautiful reverberation appropriate for the music at hand, which adds fullness and a sense of space. Consequently, a number of concert halls have built reputations for having fine acoustics, based on the quality of the perceived reverberation.

The importance of reverberation in recorded music has resulted in the the creation of artificial reverberators, electro-acoustic devices that simulate the reverberation of rooms. Early devices used springs or steel plates equipped with transducers. The advent of digital electronics has replaced these devices with the modern digital reverberator, which simulates reverberation using a linear discrete-time filter. These devices are ubiquitous in the audio production industry. Almost every bit of audio that we hear from recordings, radio, television, and movies has had artificial reverberation added. Artificial reverberation has recently found another application in the field of virtual environments, where simulating room acoustics is critical for producing a convincing immersive experience.

The subject of this paper is the study of signal processing algorithms that simulate natural room reverberation. The emphasis will be on efficient algorithms that can be implemented for real-time performance.

3.1.1 Reverberation as a linear filter

From a signal processing standpoint, it is convenient to think of a room containing sound sources and listeners as a system with inputs and outputs, where the input and output signal amplitudes correspond to acoustic variables at points in the room. For example, consider a system with one input associated with a spherical sound source, and two outputs associated with the acoustical pressures at the eardrums of a listener. To the extent that the room can be considered a linear, time-invariant (LTI) system¹, a stereo transfer function completely describes the transformation of sound pressure from the source to the ears of a listener. We can therefore simulate the effect of the room by convolving an input signal with the *binaural impulse response* (BIR):

$$\begin{aligned} y_L(t) &= \int_0^{\infty} h_L(\tau)x(t - \tau)d\tau \\ y_R(t) &= \int_0^{\infty} h_R(\tau)x(t - \tau)d\tau \end{aligned} \quad (3.1)$$

where $h_L(t)$ and $h_R(t)$ are the system impulse responses for the left and right ear, respectively; $x(t)$ is the source sound; and $y_L(t)$ and $y_R(t)$ are the resulting signals

for the left and right ear, respectively. This concept is easily generalized to the case of multiple sources and multiple listeners.

3.1.2 Approaches to reverberation algorithms

We will speak of a *reverberation algorithm*, or more simply, a *reverberator*, as a linear discrete-time system that simulates the input-output behavior of a real or imagined room. The problem of designing a reverberator can be approached from a physical or perceptual point of view.

The physical approach. The physical approach seeks to simulate exactly the propagation of sound from the source to the listener for a given room. The preceding discussion of binaural impulse responses suggests an obvious way to do this, by simply measuring the binaural impulse response of an existing room, and then rendering the reverberation by convolution.

When the room to be simulated doesn't exist, we can attempt to predict its impulse response based on purely physical considerations. This requires detailed knowledge of the geometry of the room, properties of all surfaces in the room, and the positions and directivities of the sources and receivers. Given this prior information, it is possible to apply the laws of acoustics regarding wave propagation and interaction with surfaces to predict how the sound will propagate in the space. This technique has been termed *auralization* in the literature and is an active area of research [Kleiner et al., 1993]. Typically, an auralization system first computes the impulse response of the specified room, for each source-receiver pair. These finite impulse response (FIR) filters are then used to render the room reverberation.

The advantage of this approach is that it offers a direct relation between the physical specification of the room and the resulting reverberation. However, this approach is computationally expensive and rather inflexible. Compared to other algorithms we will study, real-time convolution with a large filter response is somewhat expensive, even using an efficient algorithm. Furthermore, there is no easy way to achieve real-time parametric control of the perceptual characteristics of the resulting reverberation without recalculating a large number of FIR filter coefficients.

The perceptual approach. The perceptual approach seeks to reproduce only the perceptually salient characteristics of reverberation. Let us assume that the space of all percepts caused by reverberation can be spanned by N independent dimensions, which correspond to independently perceivable attributes of reverberation. If each perceptual attribute can be associated with a physical feature of the impulse response, then we can attempt to construct a digital filter with N parameters that reproduces exactly these N attributes. In order to simulate the reverberation from a particular room, we can measure the room response, estimate the N parameters by analyzing the impulse

response, and then plug the parameter estimates into our "universal" reverberator. The reverberator should then produce reverberation that is indistinguishable from the original, even though the fine details of the impulse responses may differ considerably.

This approach has many potential advantages:

- The reverberation algorithm can be based on efficient infinite impulse response (IIR) filters.
- The reverberation algorithm will provide real-time control of all the perceptually relevant parameters. The parameters do not need to be correlated as they often are in real rooms.
- Ideally, only one algorithm is required to simulate all reverberation.
- Existing rooms can be simulated using the analysis/synthesis approach outlined above.

One disadvantage of this method is that it doesn't necessarily provide an easy way to change a physical property of the simulated room.

The perceptually motivated method is essentially the approach that has been taken in the design of reverberation algorithms, with several caveats. First, there is a great deal of disagreement as to what the perceivable attributes of reverberation are, and how to measure these from an impulse response. Second, it is difficult to design digital filters to reproduce these attributes. Consequently, the emphasis has been to design reverberators that are *perceptually indistinguishable* from real rooms, without necessarily providing the reverberator with a complete set of independent perceptual controls.

In this paper, we will concentrate on the perceptually motivated method, because the resulting recursive algorithms are more practical and useful. We first present a concise physical and perceptual background for our study of reverberation, then discuss algorithms to simulate early reverberation, and conclude with a discussion of late reverberation algorithms.

3.2 PHYSICAL AND PERCEPTUAL BACKGROUND

The process of reverberation starts with the production of sound at a location within a room. The acoustic pressure wave expands radially outward, reaching walls and other surfaces where energy is both absorbed and reflected. Technically speaking, all reflected energy is *reverberation*. Reflection off large, uniform, rigid surfaces produces a reflection the way a mirror reflects light, but reflection off non-uniform surfaces is a complicated process, generally leading to a diffusion of the sound in various directions. The wave propagation continues indefinitely, but for practical purposes we can consider the propagation to end when the intensity of the wavefront falls below the intensity of the ambient noise level.

Assuming a direct path exists between the source and the listener, the listener will first hear the *direct sound*, followed by reflections of the sound off nearby surfaces, which are called *early echoes*. After a few hundred milliseconds, the number of reflected waves becomes very large, and the remainder of the reverberant decay is characterized by a dense collection of echoes traveling in all directions, whose intensity is relatively independent of location within the room. This is called *late reverberation* or *diffuse reverberation*, because there is equal energy propagating in all directions. In a perfectly diffuse soundfield, the energy lost due to surface absorption is proportional to the energy density of the soundfield, and thus diffuse reverberation decays exponentially with time. The time required for the reverberation level to decay to 60 dB below the initial level is defined as the *reverberation time*.

3.2.1 Measurement of reverberation

Measuring reverberation in a room usually consists of measuring an impulse response for a specific source and receiver. Pistol shots, balloon pops, and spark generators can be used as impulsive sources. Another possibility is to use an omnidirectional speaker driven by an electronic signal generator. Typical measurement signals include clicks, chirps (also known as time delay spectrometry [Heyser, 1967]), and various pseudo-random noise signals, such as maximum length (ML) sequences [Rife and Vanderkooy, 1987] and Golay codes [Foster, 1986]. The click (unit impulse) signal allows a direct measurement of the impulse response, but results in poor signal to noise ratio (SNR) because the signal energy is small for a given peak amplitude. The chirp and noise signals have significantly greater energy for a given peak amplitude, and allow the impulse response to be measured with improved SNR by deconvolving the impulse response from the recorded signal. The measurement signals are deliberately chosen to make the deconvolution easy to perform.

Figure 3.1 shows the impulse response of a concrete stairwell, plotting pressure as a function of time. The direct response is visible at the far left, followed by some early echoes, followed by the exponentially decaying late reverberation. The early echoes have greater amplitude than the direct response due to the directivities of the measurement speaker and microphone.

Rooms may contain a large number of sources with different positions and directivity patterns, each producing an independent signal. The reverberation created in a concert hall by a symphony orchestra cannot be characterized by a single impulse response. Fortunately, the statistical properties of late reverberation do not change significantly as a function of position. Thus, a point to point impulse response does characterize the late reverberation of the room, although the early echo pattern is dependent on the positions and directivities of the source and receiver.




Figure 3.1 Impulse response of reverberant stairwell measured using ML sequences.

The fact that the early and late reverberation have different physical and perceptual properties permits us to logically split the study of reverberation into early and late reverberation.

3.2.2 Early reverberation

Early reverberation is most easily studied by considering a simple geometrical model of the room. These models depend on the assumption that the dimensions of reflective surfaces in the room are large compared to the wavelength of the sound. Consequently, the sound wave may be modeled as a ray that is normal to the surface of the wavefront and reflects specularly, like light bouncing off a mirror, when the ray encounters a wall surface. Figure 3.2 shows a wall reflection using the ray model. The source is at point A , and we are interested in how sound will propagate to a listener at point B .

The reflected ray may also be constructed by considering the mirror image of the source as reflected across the plane of the wall. In figure 3.2, the *image source* thus constructed is denoted A' . This technique of reflecting sources across wall surfaces is called the *source image method*. The method allows a source with reflective boundaries to be modeled as multiple sources with no boundaries.

The image source A' is a first order source, corresponding to a sound path with a single reflection. Higher order sources corresponding to sound paths with multiple




Figure 3.2 Single wall reflection and corresponding image source A' .

reflections are created by reflecting lower order sources across wall boundaries. Frequently the resulting sources are “invisible” to the listener position, and this condition must be tested explicitly for each source. When the room is rectangular, as shown in figure 3.3, the pattern of image sources is regular and trivial to calculate. Calculation of the image source positions in irregularly-shaped rooms is more difficult, but the problem has been solved in detail [Borish, 1984]. The number of image sources of order k is roughly N^k , where N is the number of wall surfaces. The source image method is impractical for studying late reverberation because the number of sources increases exponentially, and the simplified reflection model becomes inaccurate.




Figure 3.3 A regular pattern of image sources occurs in an ideal rectangular room.

In order to calculate the impulse response at the listener’s position, the contributions from all sources are summed. Each source contributes a delayed impulse (echo), whose time delay is equal to the distance between the source and the listener divided by the speed of sound. The echo amplitude is inversely proportional to the distance

travelled, to account for spherical expansion of the sound, and proportional to the product of the reflection coefficients of the surfaces encountered. This model ignores any frequency dependent absorption, which normally occurs during surface reflections and air propagation. A more accurate model uses linear filters to approximate these frequency dependent losses [Lehnert and Blauert, 1992], such that the spectrum of each echo reaching the listener is determined by the product of the transfer functions involved in the history of that echo:

$$A(\omega) = G(\omega) \prod_{j \in S} \Gamma_j(\omega) \quad (3.2)$$

where $A(\omega)$ is the spectrum of the echo, S is the set of walls encountered, $\Gamma_j(\omega)$ is the frequency dependent transfer function that models reflection with the j th wall, and $G(\omega)$ models the absorptive losses and time delay due to air propagation.

The simplifying assumptions that permit us to consider only specular reflections are no longer met when the wall surfaces contain features that are comparable in size to the wavelength of the sound. In this case, the reflected sound will be scattered in various directions, a phenomenon referred to as *diffusion*. The source image model cannot be easily extended to handle diffusion. Most auralization systems use another geometrical model, called *ray tracing* [Krokstad et al., 1968], to model diffuse reflections. A discussion of these techniques is beyond the scope of this paper.

The early response consists largely of discrete reflections that come from specific directions, and we now consider how to reproduce the directional information. It is well known that the auditory cues for sound localization are embodied in the transformation of sound pressure by the torso, head, and external ear (pinna) [Blauert, 1983]. A *head-related transfer function* (HRTF) is a frequency response that describes this transformation from a specific free field source position to the eardrum. HRTFs are usually measured using human subjects or dummy-head microphones, and consist of response pairs, for the left and right ears, corresponding to a large number of source positions surrounding the head. When computing the binaural transfer function of a room using the geometrical models just discussed, we must convolve each directional echo with the HRTF corresponding to the direction of the echo [Wightman and Kistler, 1989, Begault, 1994].

The binaural directional cues captured by HRTFs are primarily the interaural time difference (ITD) and interaural intensity difference (IID) which vary as a function of frequency. Echoes that arrive from lateral directions (i.e. from either side of the listener) are important for modifying the spatial character of the perceived reverberation. The ITD of a lateral sound source is well modeled by a delay corresponding to the difference in path lengths between the two ears. Similarly, the IID may be modeled as a lowpass filtering of the signal arriving at the opposite (contralateral) ear.




Figure 3.15 Mixing matrix \mathbf{M} used to form uncorrelated outputs from parallel comb filters [Schroeder, 1962]. $A_i(z)$ are allpass filters, and $C_i(z)$ are comb filters.

$$\theta = \arcsin(IACC)/2$$




Figure 3.16 Controlling IACC in binaural reverberation [Martin et al., 1993, Jot, 1992b].

3.4.5 Moorer's reverberator

Schroeder's original reverberator sounds quite good, particularly for short reverberation times and moderate reverberation levels. For longer reverberation times or higher levels, some sonic deficiencies become noticeable and these have been described by various authors [Moorer, 1979, Griesinger, 1989, Jot and Chaigne, 1991]:

- The initial response sounds too discrete, leading to a grainy sound quality, particularly for impulsive input sounds, such as a snare drum.
- The amplitude of the late response, rather than decaying smoothly, can exhibit unnatural modulation, often described as a fluttering or beating sound.
- For longer reverberation times, the reverberation sounds tonally colored, usually referred to as a metallic timbre.

- The echo density is insufficient, and doesn't increase with time.

All reverberation algorithms are susceptible to one or more of these faults, which usually do not occur in real rooms, certainly not good sounding ones. In addition to these criticisms, there is the additional problem that Schroeder's original proposal does not provide a frequency dependent reverberation time.

Moorer later reconsidered Schroeder's reverberator and made several improvements [Moorer, 1979]. The first of these was to increase the number of comb filters from 4 to 6. This was necessary in order to effect longer reverberation times, while maintaining sufficient frequency and echo density according to equation 3.27. Moorer also inserted a one-pole lowpass filter into each comb filter feedback loop, as shown in figure 3.17. The cutoff frequencies of the lowpass filters were based on a physical consideration of the absorption of sound by air. Adding the lowpass filters caused the reverberation time to decrease at higher frequencies and Moorer noted that this made the reverberation sound more realistic. In addition, several other benefits were observed. The response to impulsive sounds was greatly improved, owing to the fact that the impulses are smoothed by the lowpass filtering. This improves the subjective quality of both the early response and the late response, which suffers less from a metallic sound quality or a fluttery decay.




Figure 3.17 Comb filter with lowpass filter in feedback loop [Moorer, 1979].

Despite these improvements many problems remained. The frequency dependent reverberation time is the net result of the lowpass filtering, but it is not possible to specify a function $T_r(\omega)$ which defines the reverberation time as a function of frequency. Furthermore, the recurring problems of metallic sounding decay and fluttery late response are reduced but not entirely eliminated by this reverberator.

3.4.6 Allpass reverberators

We now study reverberators that are based on a series association of allpass filters². Schroeder experimented with reverberators consisting of 5 allpass filters in series, with delays starting at 100 msec and decreasing roughly by factors of 1/3, and with gains of about 0.7 [Schroeder, 1962]. Schroeder noted that these reverberators were indistinguishable from real rooms in terms of coloration, which may be true with

stationary input signals, but other authors have found that series allpass filters are extremely susceptible to tonal coloration, especially with impulsive inputs [Moorer, 1979, Gardner, 1992]. Moorer experimented with series allpass reverberators, and made the following comments [Moorer, 1979]:

- The higher the order of the system, the longer it takes for the echo density to build up to a pleasing level.
- The smoothness of the decay depends critically on the particular choice of the delay and gain parameters.
- The decay exhibits an annoying, metallic ringing sound.

The z transform of a series connection of N allpass filters is:

$$H(z) = \prod_{i=1}^N \frac{z^{-m_i} - g_i}{1 - g_i z^{-m_i}} \quad (3.30)$$

where m_i and g_i are the delay and gain, respectively, of allpass filter i . It is possible to ensure that the pole moduli are all the same, by basing the gains on the delay length as indicated by equation 3.19. However, this does not solve the problem of the metallic sounding decay.

Gardner has described reverberators based on a "nested" allpass filter, where the delay of an allpass filter is replaced by a series connection of a delay and another allpass filter [Gardner, 1992]. This type of allpass filter is identical to the lattice form shown in figure 3.18. Several authors have suggested using nested allpass filters for reverberators [Schroeder, 1962, Gerzon, 1972, Moorer, 1979]. The general form of such a filter is shown in figure 3.19, where the allpass delay is replaced with a system function $A(z)$, which is allpass. The transfer function of this form is written:

$$H(z) = \frac{A(z) - g}{1 - gA(z)} \quad (3.31)$$

The magnitude squared response of $H(z)$ is:

$$|H(z)|^2 = \frac{|A(z)|^2 - 2g\text{Re}\{A(z)\} + g^2}{1 - 2g\text{Re}\{A(z)\} + g^2|A(z)|^2} = 1 \text{ if } |A(z)| = 1 \quad (3.32)$$

which is verified to be allpass if $A(z)$ is allpass [Gardner, 1992, Jot, 1992b]. This filter is not realizable unless $A(z)$ can be factored into a delay in series with an allpass filter, otherwise a closed loop is formed without delay. The advantage of using a nested allpass structure can be seen in the time domain. Echoes created by the inner allpass filter are recirculated to itself via the outer feedback path. Thus, the echo density of a nested allpass filter increases with time, as in real rooms.




Figure 3.18 Lattice allpass structure.




Figure 3.19 Generalization of figure 3.18.

A useful property of allpass filters is that no matter how many are nested or cascaded in series, the response is still allpass. This makes it very easy to verify the stability of the resulting system, regardless of complexity. Gardner suggested a general structure for a monophonic reverberator constructed with allpass filters, shown in figure 3.20 [Gardner, 1992]. The input signal flows through a cascade of allpass sections $A_i(z)$, and is then recirculated upon itself through a lowpass filter $H_{LP}(z)$ and an attenuating gain g . Gardner noted that when the output of the allpass filters was recirculated to the input through a sufficient delay, the characteristic metallic sound of the series allpass was greatly reduced.

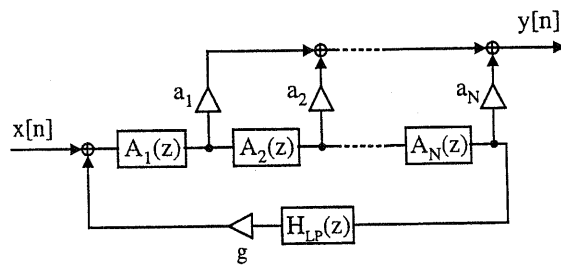


Figure 3.20 Reverberator formed by adding absorptive losses to an allpass feedback loop [Gardner, 1992].

The output is formed as a linear combination of the outputs of the allpass sections. The stability of the system is guaranteed, provided the magnitude of the loop gain is less than 1 for all frequencies (i.e. $|gH_{LP}(e^{j\omega})| < 1$ for all ω). The overall transfer function of this system is in general not allpass, due to phase cancellation between the output taps and also the presence of the outer feedback loop. As the input signal is diffused through the allpass filters, each tap outputs a different response shape. Consequently, it is possible to customize the amplitude envelope of the reverberant decay by adjusting the coefficients a_i . The reverberation time can be adjusted by changing the feedback gain g . The lowpass filter simulates frequency dependent absorptive losses, and lower cutoff frequencies generally result in a less metallic sounding, but duller, late response.

Figure 3.21 shows a complete schematic of an allpass feedback loop reverberator described by Dattorro [Dattorro, 1997], who attributes this style of reverberator to Griesinger. The circuit is intended to simulate an electro-acoustical plate reverberator, characterized by a rapid buildup of echo density followed by an exponential reverberant decay. The monophonic input signal passes through several short allpass filters, and then enters what Dattorro terms the reverberator "tank", consisting of two systems like that of figure 3.20 which have been cross-coupled. This is a useful structure for producing uncorrelated stereo outputs, which are obtained by forming weighted sums of taps within the tank. The reverberator incorporates a time varying delay element in each of the cross-coupled systems. The purpose of the time varying delays is to further decrease tonal coloration by dynamically altering the resonant frequencies.

There are many possible reverberation algorithms that can be constructed by adding absorptive losses to allpass feedback loops, and these reverberators can sound very good. However, the design of these reverberators has to date been entirely empirical. There is no way to specify in advance a particular reverberation time function $T_r(\omega)$, nor is there a deterministic method for choosing the filter parameters to eliminate tonal coloration.

3.5 FEEDBACK DELAY NETWORKS

Gerzon generalized the notion of unitary multichannel networks, which are N -dimensional analogues of the allpass filter [Gerzon, 1976]. An N -input, N -output LTI system is defined to be *unitary* if it preserves the total energy of all possible input signals. Similarly, a matrix \mathbf{M} is unitary if $\|\mathbf{M}\mathbf{u}\| = \|\mathbf{u}\|$ for all vectors \mathbf{u} , which is equivalent to requiring that $\mathbf{M}^T\mathbf{M} = \mathbf{M}\mathbf{M}^T = \mathbf{I}$, where \mathbf{I} is the identity matrix. It is trivial to show that the product of two unitary matrices is also unitary, and consequently the series cascade of two unitary systems is a unitary system. Simple unitary systems we have encountered include a set of N delay lines, and a set of N allpass filters. It is also easy to show that an N -channel unitary system and an M -channel unitary system can be combined to form an $(N + M)$ channel unitary system by diagonally juxtaposing their system matrices. Gerzon showed that a feedback modification can be




Figure 3.21 Dattorro's plate reverberator based on an allpass feedback loop, intended for 29.8 kHz sampling rate [Dattorro, 1997]. $H_1(z)$ and $H_2(z)$ are low-pass filters described in figure 3.11; $H_1(z)$ controls the bandwidth of signals entering the reverberator, and $H_2(z)$ controls the frequency dependent decay. Stereo outputs y_L and y_R are formed from taps taken from labelled delays as follows: $y_L = a[266] + a[2974] - b[1913] + c[1996] - d[1990] - e[187] - f[1066]$, $y_R = d[353] + d[3627] - e[1228] + f[2673] - a[2111] - b[335] - c[121]$. In practice, the input is also mixed with each output to achieve a desired reverberation level. The time varying functions $u(t)$ and $v(t)$ are low frequency (≈ 1 Hz) sinusoids that span 16 samples peak to peak. Typical coefficients values are $g_1 = 0.75$, $g_2 = 0.625$, $g_3 = 0.7$, $g_4 = 0.5$, $g_5 = 0.9$.

made to a unitary system without destroying the unitary property [Gerzon, 1976], in a form completely analogous to the feedback around the delay in an allpass filter. Gerzon applied these principles to the design of multichannel reverberators, and suggested the basic feedback topologies found in later work [Gerzon, 1971, Gerzon, 1972].

Stautner and Puckette proposed a four channel reverberator consisting of four delay lines with a feedback matrix [Stautner and Puckette, 1982], shown in figure 3.22. The feedback matrix allows the output of each delay to be recirculated to each delay input, with the matrix coefficients controlling the weights of these feedback paths. The structure can be seen as a generalization of Schroeder's parallel comb filter, which would arise using a diagonal feedback matrix. This structure is capable of much higher echo densities than the parallel comb filter, given a sufficient number of non-zero feedback coefficients and incommensurate delay lengths. The delays were chosen in accordance with Schroeder's suggestions.

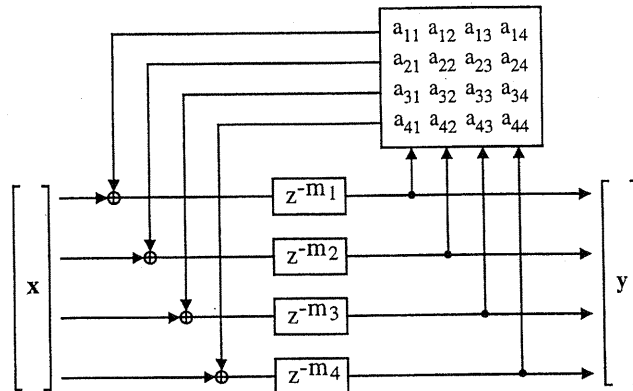


Figure 3.22 Stautner and Puckette's four channel feedback delay network [Stautner and Puckette, 1982].

Stautner and Puckette make a number of important points regarding this system:

- Stability is guaranteed if the feedback matrix A is chosen to be the product of a unitary matrix and a gain coefficient g , where $|g| < 1$. They suggest the matrix:

$$A = g \frac{1}{\sqrt{2}} \begin{bmatrix} 0 & 1 & 1 & 0 \\ -1 & 0 & 0 & -1 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & -1 & 0 \end{bmatrix} \quad (3.33)$$

where g controls the reverberation time. If $|g| = 1$, A is unitary.

- The outputs will be mutually incoherent, and thus can be used in a four channel loudspeaker system to render a diffuse soundfield.
- Absorptive losses can be simulated by placing a lowpass filter in series with each delay line.
- The early reverberant response can be customized by injecting the input signal appropriately into the interior of the delay lines.

The authors note that fluttering and tonal coloration is present in the late decay of this reverberator. They attribute the fluttering to the beating of adjacent modes, and suggest that the beat period be made greater than the reverberation time by suitably reducing the mean spacing of modes according to equation 3.25. To reduce the tonal coloration, they suggest randomly varying the lengths of the delays.

3.5.1 Jot's reverberator

We now discuss the recent and important work by Jot, who has proposed a reverberator structure with two important properties [Jot, 1992b]:

- A reverberator can be designed with arbitrary time and frequency density while simultaneously guaranteeing absence of tonal coloration in the late decay.
- The resulting reverberator can be specified in terms of the desired reverberation time $T_r(\omega)$ and frequency response envelope $G(\omega)$.

This is accomplished by starting with an energy conserving system whose impulse response is perceptually equivalent to stationary white noise. Jot calls this a *reference filter*, but we will also use the term *lossless prototype*. Jot chooses lossless prototypes from the class of unitary feedback systems. In order to effect a frequency dependent reverberation time, absorptive filters are associated with each delay in the system. This is done in a way that eliminates coloration in the late response, by guaranteeing the local uniformity of pole modulus.

Jot generalizes the notion of a monophonic reverberator using the *feedback delay network* (FDN) structure shown in figure 3.23. The structure is a completely general specification of a linear system containing N delays.

Using vector notation and the z transform, the equations for the output of the system $y(z)$ and the delay lines $s_i(z)$ are [Jot and Chaigne, 1991]:

$$y(z) = \mathbf{c}^T \mathbf{s}(z) + dx(z) \quad (3.34)$$

$$\mathbf{s}(z) = \mathbf{D}(z)[\mathbf{A}\mathbf{s}(z) + \mathbf{b}x(z)] \quad (3.35)$$

where:

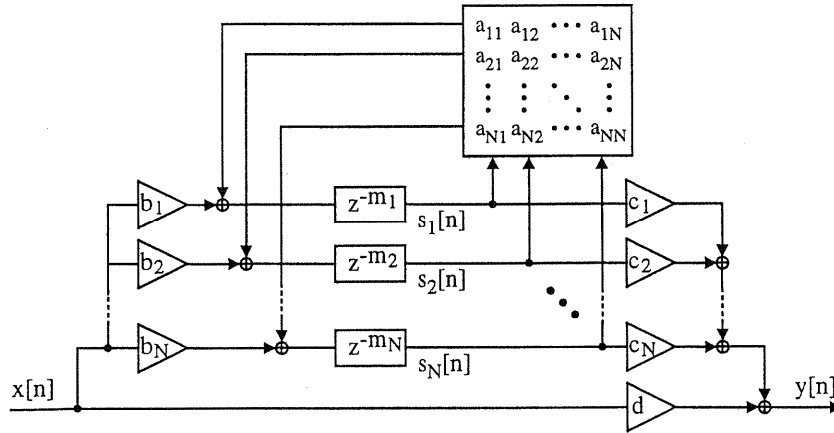


Figure 3.23 Feedback delay network as a general specification of a reverberator containing N delays [Jot and Chaigne, 1991]

$$s(z) = \begin{bmatrix} s_1(z) \\ \vdots \\ s_N(z) \end{bmatrix} \quad \mathbf{b} = \begin{bmatrix} b_1 \\ \vdots \\ b_N \end{bmatrix} \quad \mathbf{c} = \begin{bmatrix} c_1 \\ \vdots \\ c_N \end{bmatrix} \quad (3.36)$$

$$\mathbf{D}(z) = \begin{bmatrix} z^{-m_1} & & 0 \\ & \ddots & \\ 0 & & z^{-m_N} \end{bmatrix} \quad \mathbf{A} = \begin{bmatrix} a_{11} & \cdots & a_{1N} \\ \vdots & \ddots & \vdots \\ a_{N1} & \cdots & a_{NN} \end{bmatrix} \quad (3.37)$$

The FDN can be extended to multiple inputs and outputs by replacing the vectors \mathbf{b} and \mathbf{c} with appropriate matrices. The system transfer function is obtained by eliminating $s(z)$ from the preceding equations [Jot and Chaigne, 1991]:

$$H(z) = \frac{y(z)}{x(z)} = \mathbf{c}^T [\mathbf{D}(z^{-1}) - \mathbf{A}]^{-1} \mathbf{b} + d \quad (3.38)$$

The system zeros are given by [Rocchesso and Smith, 1994]:

$$\det[\mathbf{A} - \frac{\mathbf{bc}^T}{d} - \mathbf{D}(z^{-1})] = 0 \quad (3.39)$$

The system poles are given by those values of z that nullify the denominator of equation 3.38, in other words the solutions to the characteristic equation:

$$\det[\mathbf{A} - \mathbf{D}(z^{-1})] = 0 \quad (3.40)$$

Assuming \mathbf{A} is a real matrix, the solutions to the characteristic equation 3.40 will either be real or complex-conjugate pole pairs. Equation 3.40 is not easy to solve in the general case, but for specific choices of \mathbf{A} the solution is straightforward. For instance, when \mathbf{A} is diagonal, the system represents Schroeder's parallel comb filter, and the poles are given by equation 3.21. More generally, when \mathbf{A} is triangular, the matrix $\mathbf{A} - \mathbf{D}(z^{-1})$ is also triangular; and because the determinant of a triangular matrix is the product of the diagonal entries, equation 3.40 reduces to:

$$\prod_{i=1}^N (a_{ii} - z^{m_i}) = 0 \quad (3.41)$$

This is verified to be identical to equation 3.21. Any series combination of elementary filters – for instance, a series allpass filter – can be expressed as a feedback delay network with a triangular feedback matrix [Jot and Chaigne, 1991].

3.5.2 Unitary feedback loops

Another situation that interests us occurs when the feedback matrix \mathbf{A} is chosen to be unitary, as suggested by Stautner and Puckette. Because the set of delay lines is also a unitary system, a unitary feedback loop is formed by the cascade of the two unitary systems. A general form of this situation is shown in figure 3.24, where $U_1(z)$ corresponds to the delay matrix, and $U_2(z)$ corresponds to the feedback matrix.




Figure 3.24 Unitary feedback loop [Jot, 1992b].

Because a unitary system preserves the energy of input signals, it is intuitively obvious that a unitary feedback loop will conserve energy. It can be shown that the system poles of a unitary feedback loop all have unit modulus, and thus the system response consists of non-decaying eigenmodes [Jot, 1992b].

Another way to demonstrate this is to consider the state variable description for the FDN shown in figure 3.23. It is straightforward to show that the resulting state transition matrix is unitary if and only if the feedback matrix \mathbf{A} is unitary [Jot, 1992b, Rocchesso and Smith, 1997]. Thus, a unitary feedback matrix is sufficient to create a lossless

FDN prototype. However, we will later see that there are other choices for the feedback matrix that also yield a lossless system.

3.5.3 Absorptive delays

Jot has demonstrated that unitary feedback loops can be used to create lossless prototypes whose impulse responses are perceptually indistinguishable from stationary white noise [Jot and Chaigne, 1991]. Moorer previously noted that convolving source signals with exponentially decaying Gaussian white noise produces a very natural sounding reverberation [Moorer, 1979]. Consequently, by introducing absorptive losses into a suitable lossless prototype, we should obtain a natural sounding reverberator. Jot's method for introducing absorptive losses guarantees that the colorless quality of the lossless prototype is maintained. This is accomplished by associating a gain $k_i < 1$ with each delay i in the filter, as shown in figure 3.25.




Figure 3.25 Associating an attenuation with a delay.

The logarithm of the the gain is proportional to the length of the delay:

$$k_i = \gamma^{m_i} \quad (3.42)$$

Provided all the delays are so modified, this has the effect of replacing z with z/γ in the expression for the system function $H(z)$, regardless of the filter structure. Starting from a lossless prototype whose poles are all on the unit circle, the above modification will cause all the poles to have a modulus equal to γ . Therefore, the lossless prototype response $h[n]$ will be multiplied by an exponential envelope γ^n where γ is the decay factor per sampling period [Jot and Chaigne, 1991, Jot, 1992b]. By maintaining the uniformity of pole modulus, we avoid the situation where the response in the neighborhood of a frequency is dominated by a few poles with relatively large moduli.

The decay envelope is made frequency dependent by specifying frequency dependent losses in terms of the reverberation time $T_r(\omega)$. This is accomplished by associating with each delay i an absorptive filter $h_i(z)$, as shown in figure 3.26. The filter is chosen such that the logarithm of its magnitude response is proportional to the delay length and inversely proportional to the reverberation time, as suggested by equation 3.19 [Jot and Chaigne, 1991]:

$$20 \log_{10} |h_i(e^{j\omega})| = \frac{-60T}{T_r(\omega)} m_i \quad (3.43)$$

This expression ignores the phase response of the absorptive filter, which has the effect




Figure 3.26 Associating an absorptive filter with a delay.

of slightly modifying the effective length of the delay. In practice, it is not necessary to take the phase delay into consideration [Jot and Chaigne, 1991]. By replacing each delay with an absorptive delay as described above, the poles of the prototype filter no longer appear on a circle centered at the origin, but now lie on a curve specified by the reverberation time $T_r(\omega)$.

A consequence of incorporating the absorptive filters into the lossless prototype is that the frequency response envelope of the reverberator will no longer be flat. For exponentially decaying reverberation, the frequency response envelope is proportional to the reverberation time at all frequencies. We can compensate for this effect by associating a correction filter $t(z)$ in series with the reference filter, whose squared magnitude is inversely proportional to the reverberation time [Jot, 1992b]:

$$|t(e^{j\omega})| \propto \frac{1}{\sqrt{T_r(\omega)}} \quad (3.44)$$

After applying the correction filter, the frequency response envelope of the reverberator will be flat. This effectively decouples the reverberation time control from the overall gain of the reverberator. The final reverberator structure is shown in figure 3.27. Any additional equalization of the reverberant response, for instance, to match the frequency envelope of an existing room, can be effected by another filter in series with the correction filter.

3.5.4 Waveguide reverberators

Smith has proposed multichannel reverberators based on a *digital waveguide network* (DWN) [Smith, 1985]. Each waveguide is a bi-directional delay line, and junctions between multiple waveguides produce lossless signal scattering. Figure 3.28 shows an N -branch DWN which is isomorphic to the N -delay FDN shown in figure 3.23 [Smith and Rocchesso, 1994].

The waves travelling into the junction are associated with the FDN delay line outputs $s_i[n]$. The length of each waveguide is half the length of the corresponding FDN delay, because the waveguide signal must make a complete round trip to return to the scattering junction. An odd-length delay can be accommodated by replacing the non-inverting reflection with a unit sample delay.

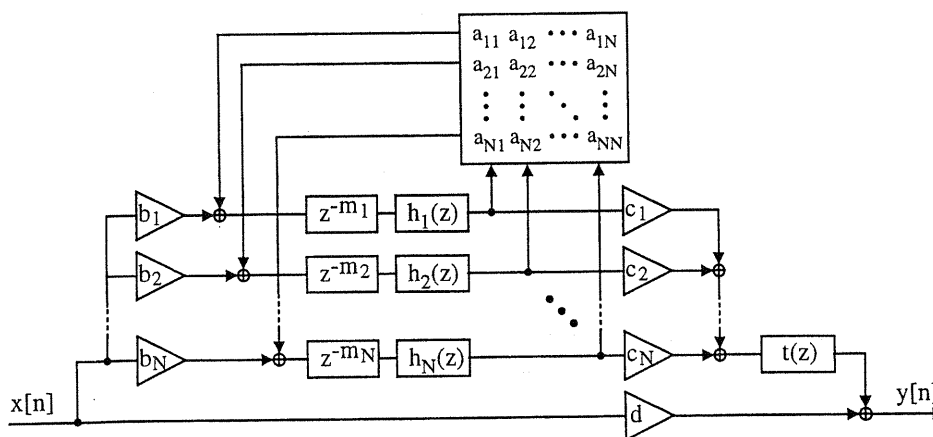


Figure 3.27 Reverberator constructed by associating a frequency dependent absorptive filter with each delay of a lossless FDN prototype filter [Jot and Chaigne, 1991].




Figure 3.28 Waveguide network consisting of a single scattering junction to which N waveguides are attached. Each waveguide is terminated by an ideal non-inverting reflection, indicated by a black dot [Smith and Rocchesso, 1994].

The usual DWN notation defines the incoming and outgoing pressure variables as $p_i^+ = s_i[n]$ and $p_i^- = s_i[n + m_i]$, respectively, and therefore the operation of the scattering junction can be written in vector notation as

$$\mathbf{p}^- = \mathbf{A}\mathbf{p}^+ \quad (3.45)$$

where \mathbf{A} is interpreted as a *scattering matrix* associated with the junction.

As we have already discussed, a lossless FDN results when the feedback matrix is chosen to be unitary. Smith and Rocchesso have shown that the waveguide interpretation leads to a more general class of lossless scattering matrices [Smith and Rocchesso, 1994]. This is due to the fact that each waveguide may have a different characteristic admittance. A scattering matrix is lossless if and only if the active complex power is scattering-invariant, i.e., if and only if

$$\begin{aligned} \mathbf{p}^{+*} \Gamma \mathbf{p}^+ &= \mathbf{p}^{-*} \Gamma \mathbf{p}^- \\ \Rightarrow \mathbf{A}^* \Gamma \mathbf{A} &= \Gamma \end{aligned}$$

where Γ is a Hermitian, positive-definite matrix which can be interpreted as a generalized junction admittance. For the waveguide in figure 3.28, we have $\Gamma = \text{diag}(\Gamma_1, \dots, \Gamma_N)$, where Γ_i is the characteristic admittance of waveguide i . When \mathbf{A} is unitary, we have $\Gamma = \mathbf{I}$. Thus, unitary feedback matrices correspond to DWNs where the waveguides all have unit characteristic admittance, or where the signal values are in units of root power [Smith and Rocchesso, 1994].

Smith and Rocchesso have shown that a DWN scattering matrix (or a FDN feedback matrix) is lossless if and only if its eigenvalues have unit modulus and its eigenvectors are linearly independent. Therefore, lossless scattering matrices may be fully parameterized as

$$\mathbf{A} = \mathbf{T}^{-1} \mathbf{D} \mathbf{T} \quad (3.46)$$

where \mathbf{D} is any unit modulus diagonal matrix, and \mathbf{T} is any invertible matrix [Smith and Rocchesso, 1994]. This yields a larger class of lossless scattering matrices than given by unitary matrices. However, not all lossless scattering matrices can be interpreted as a physical junction of N waveguides (e.g., consider a permutation matrix).

3.5.5 Lossless prototype structures

Jot has described many lossless FDN prototypes based on unitary feedback matrices. A particularly useful unitary feedback matrix \mathbf{A}_N , which maximizes echo density while reducing implementation cost, is taken from the class of Householder matrices [Jot, 1992b]:

$$\mathbf{A}_N = \mathbf{J}_N - \frac{2}{N} \mathbf{u}_N \mathbf{u}_N^T \quad (3.47)$$

where \mathbf{J}_N is an $N \times N$ permutation matrix, and \mathbf{u}_N is an $N \times 1$ column vector of 1's. This unitary matrix contains only two different values, both nonzero, and thus it achieves maximum echo density when used in the structure of figure 3.27. Because $\mathbf{u}_N \mathbf{u}_N^T$ is a matrix containing all 1's, computation of $\mathbf{A}_N \mathbf{x}$ consists of permuting the elements of \mathbf{x} according to \mathbf{J}_N , and adding to these the sum of the elements of \mathbf{x} times the factor $-2/N$. This requires roughly $2N$ operations as opposed to the N^2 operations normally required. When \mathbf{J}_N is the identity matrix \mathbf{I}_N , the resulting system is a modification of Schroeder's parallel comb filter which maximizes echo density as shown in figure 3.29.




Figure 3.29 Modification of Schroeder's parallel comb filter to maximize echo density [Jot, 1992b].

Jot has discovered that this structure produces a periodic parasitic echo with period equal to the sum of the delay lengths. This is a result of constructive interference between the output signals of the delays, and can be eliminated by choosing the coefficients c_i in figure 3.27 such that every other channel undergoes a phase inversion (multiplication by -1) [Jot, 1992b]. Another interesting possibility proposed by Jot is choosing \mathbf{J}_N to be a circular permutation matrix. This causes the delay lines to feed one another in series, which greatly simplifies the memory management in the final implementation.

Rocchesso and Smith have suggested using unitary circulant matrices for the feedback matrix of a FDN [Rocchesso and Smith, 1994]. Circulant matrices have the form:

$$\mathbf{A} = \begin{bmatrix} a[0] & a[1] & a[2] & \cdots & a[N-1] \\ a[N-1] & a[0] & a[1] & \cdots & a[N-2] \\ a[N-2] & a[N-1] & a[0] & \cdots & a[N-3] \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a[1] & a[2] & a[3] & \cdots & a[0] \end{bmatrix} \quad (3.48)$$

Multiplication by a circulant matrix implements circular convolution of a column vector with the first row of the matrix. A circulant matrix \mathbf{A} can be factored as shown in equation 3.46 where \mathbf{T} is the discrete Fourier transform (DFT) matrix and \mathbf{D} is a diagonal matrix whose elements are the DFT of the first row of \mathbf{A} . The diagonal elements of \mathbf{D} are the eigenvalues of \mathbf{A} . A circulant matrix is thus lossless (and unitary) when its eigenvalues (the spectrum of the first row) have unit modulus. The advantages of using a circulant matrix are that the eigenvalues can be explicitly specified, and computation of the product can be accomplished in $O(N \log(N))$ time using the Fast Fourier transform (FFT).

All of the late reverberator structures we have studied can be seen as an energy conserving system with absorptive losses inserted into the structure. When the absorptive losses are removed, the structure of the lossless prototype is revealed. This is true for Schroeder's parallel comb filter when the feedback coefficients are unity, which corresponds to a FDN feedback matrix equal to the identity matrix. The allpass feedback loop reverberator in figure 3.20 consists of a unitary feedback loop when absorptive losses are removed. Stautner and Puckette's FDN reverberator is also a unitary feedback loop when $|g| = 1$ (see equation 3.33). However, the method shown for adding the absorptive losses in these reverberators does not necessarily prevent coloration in the late decay. This can be accomplished by associating an absorptive filter with each delay in the reverberator according to equation 3.43.

The parameters of the reference structure are the number of delays N , the lengths of the delays m_i , and the feedback matrix coefficients. If a large number of inputs or outputs is desired, this can also affect the choice of the reference structure. The total length of the delays in seconds, equal to the modal density, should be greater than the density of frequency maxima for the room to be simulated. Thus, the minimum total length required is $T_r/4$, after equation 3.24. A total delay of 1 to 2 seconds is sufficient to produce a reference filter response that is perceptually indistinguishable from white noise [Jot, 1992b], which gives an upper bound on the total delay required for infinite reverberation times with broadband input signals. To improve the quality of the reverberation in response to narrowband input signals, one may wish to use a total delay at least equal to the maximum reverberation time desired, after equation 3.25. The number of delays and the lengths of the delays, along with the choice of feedback matrix, determines the buildup of echo density. These decisions must be made empirically by evaluating the quality of the reference filter response.

3.5.6 Implementation of absorptive and correction filters

Once a lossless prototype has been chosen, the absorptive filters and the correction filter need to be implemented based on a desired reverberation time curve. Jot has specified a simple solution using first order IIR filters for the absorptive filters, whose transfer functions are written [Jot, 1992b]:

$$h_i(z) = k_i \frac{1 - \beta_i}{1 - \beta_i z^{-1}} \quad (3.49)$$

Remarkably, this leads to a correction filter which is first order FIR:

$$t(z) = g \frac{1 - \beta z^{-1}}{1 - \beta} \quad (3.50)$$

The filter parameters are based on the reverberation time at zero frequency and the Nyquist frequency, notated $T_r(0)$ and $T_r(\pi)$, respectively:

$$k_i = 10^{-3\tau_i/T_r(0)}, \beta_i = 1 - \frac{2}{1 + k_i^{(1+1/\epsilon)}} \quad (3.51)$$

$$g = \sqrt{\frac{\sum \tau_i}{T_r(0)}}, \beta = \frac{1 - \sqrt{\epsilon}}{1 + \sqrt{\epsilon}}, \epsilon = \frac{T_r(\pi)}{T_r(0)}$$

The derivation of these parameters is detailed in the reference [Jot, 1992b]. The family of reverberation time curves obtained from first order filters is limited, but leads to natural sounding reverberation. Jot also describes methods for creating higher order absorption and correction filters by combining first order sections.

3.5.7 Multirate algorithms

Jot's method of incorporating absorptive filters into a lossless prototype yields a system whose poles lie on a curve specified by the reverberation time. An alternative method to obtain the same pole locus is to combine a bank of bandpass filters with a bank of comb filters, such that each comb filter processes a different frequency range. The feedback gain of each comb filter then determines the reverberation time for the corresponding frequency band.

This approach has been extended to a multirate implementation by embedding the bank of comb filters in the interior of a multirate analysis/synthesis filterbank [Zoelzer et al., 1990]. A multirate implementation reduces the memory requirements for the comb filters, and also allows the use of an efficient polyphase analysis/synthesis filterbank [Vaidyanathan, 1993].

3.5.8 Time-varying algorithms

There are several reasons why one might want to incorporate time variation into a reverberation algorithm. One motivation is to reduce coloration and fluttering in the reverberant response by varying the resonant frequencies. Another use of time variation is to reduce feedback when the reverberator is coupled to an electro-acoustical sound reinforcement system, as is the case in reverberation enhancement systems [Griesinger, 1991]. The time variation should always be implemented so as to yield a natural sounding reverberation free from audible amplitude or frequency modulations. There are several ways to add time variation to an existing algorithm:

- Modulate the lengths of the delays, e.g., as shown in figure 3.21.
- Vary the coefficients of the feedback matrix in the reference filter while maintaining the energy conserving property, or similarly vary the allpass gains of an allpass feedback loop reverberator.
- Modulate the output tap gains of an allpass feedback loop structure such as in figure 3.20, or similarly vary the mixing matrix shown in equation 3.28.

There are many ways to implement variable delay lines [Laakso et al., 1996]. A simple linear interpolator works well, but for better high frequency performance, it may be preferable to use a higher order Lagrangian interpolator. Dattorro has suggested using allpass interpolation, which is particularly suitable because the required modulation rate is low [Dattorro, 1997]. Obviously, modulating the delay length causes the signal passing through the delay to be frequency modulated. If the depth or rate of the modulation is too great, the modulation will be audible in the resulting reverberation. This is particularly easy to hear with solo piano music. The maximum detune should be restricted to a few cents, and the modulation rate should be on the order of 1 Hz.

The notion of changing the filter coefficients while maintaining an energy conserving system has been suggested by Smith [Smith, 1985], who describes the result as placing the signal in a changing lossless maze. Smith suggests that all coefficient modulation be done at sub-audio rates to avoid sideband generation, and warns of an "undulating" sound that can occur with slow modulation that is too deep.

Although many commercial reverberators use time variation to reduce tonal coloration, very little has been published on time-varying techniques. There is no theory which relates the amount and type of modulation to the reduction of tonal coloration in the late response, nor is there a way to predict whether the modulation will be noticeable. Consequently, all the time-varying methods are completely empirical in nature.

3.6 CONCLUSIONS

This paper has discussed algorithms for rendering reverberation in real-time. A straightforward method for simulating room acoustics is to sample a room impulse response and render the reverberation using convolution. Synthetic impulse responses can be created using auralization techniques. The availability of efficient, zero delay convolution algorithms make this a viable method for real-time room simulation. The drawback of this method is the lack of parameterized control over perceptually salient characteristics of the reverberation. This can be a problem when we attempt to use these systems in interactive virtual environments.

Reverberators implemented using recursive filters offer parameterized control due to the small number of filter coefficients. The problem of designing efficient, natural sounding reverberation algorithms has always been to avoid unpleasant coloration and fluttering in the decay. In many ways, Jot's work has revolutionized the state of the art, because it is now possible to design colorless reverberators without resorting to solely empirical design methods. It is possible to specify in advance the reverberation time curve of the reverberator, permitting an analysis/synthesis method for reverberator design which concentrates on reproducing the energy decay relief of the target room. Interestingly, many of the fundamental ideas can be traced back to Schroeder's original work, which is now more than thirty years old.

There are still problems to be solved. Reproducing a complicated reverberation time curve using Jot's method requires associating a high order filter with each delay in the lossless prototype, and this is expensive. It is an open question whether the constraint of uniform pole modulus necessarily requires one absorptive filter per delay line (Jean-Marc Jot, personal communication, 1994). Many of the commercially available reverberators probably use time-variation to reduce tonal coloration, yet the study of time-varying algorithms has received almost no attention in the literature. A general theory of tonal coloration in reverberation is needed to explain why certain algorithms sound good and others sound bad.

The study of reverberation has been fertile ground for many acousticians, psychologists, and electrical engineers. There is no doubt it will continue to be so in the future.

Acknowledgments

The author would like to express his sincerest gratitude to David Griesinger for making this paper possible, to Jean-Marc Jot for his inspiring work, to Jon Dattorro for his useful comments, and to the author's colleagues at the Machine Listening Group at the MIT Media Lab – in particular, Eric Scheirer, Keith Martin, and Dan Ellis – for helping to prepare this paper.

Notes

1. Rooms are very linear but they are not time-invariant due to the motion of people and air. For practical purposes we consider them to be LTI systems.
2. There are many ways to implement allpass filters [Moorer, 1979, Jot, 1992b]; two methods are shown in Figures 3.13 and 3.14.