

Artificial Reverberation and Spatialization

Julius Smith and Nelson Lee

RealSimple Project*

Center for Computer Research in Music and Acoustics (CCRMA)

Department of Music, Stanford University

Stanford, California 94305

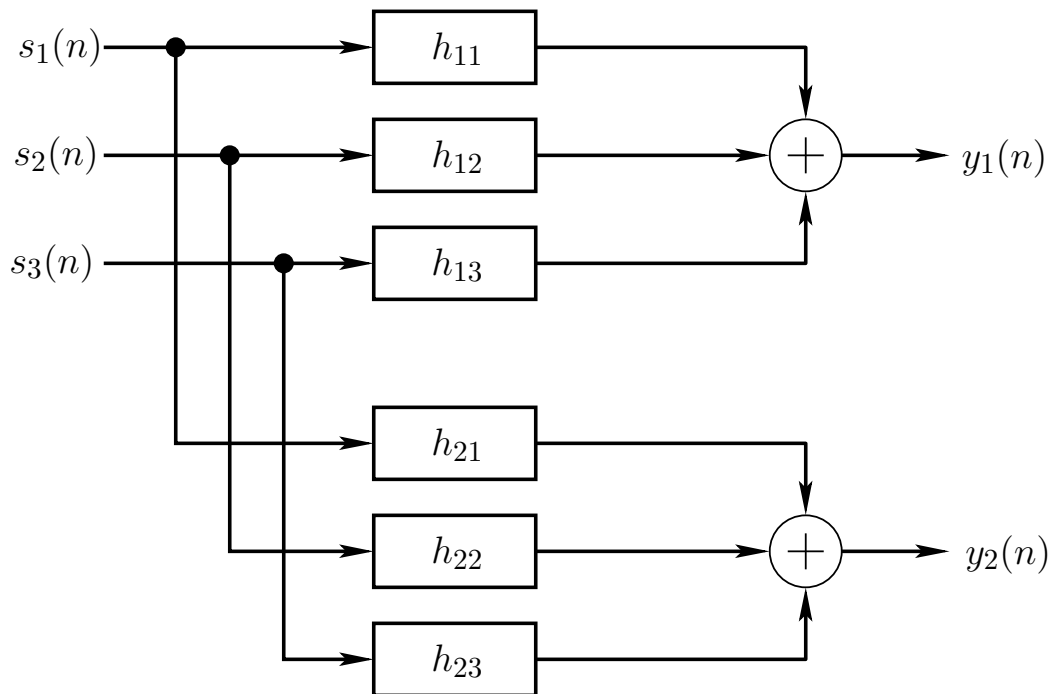
September 19, 2007

Outline

- The Reverb Problem
- Reverb Perception
- Early Reflections
- Late Reverb
- Schroeder Reverbs
- Feedback Delay Network (FDN) Reverberators
- Waveguide Reverberators

*Work supported by the Wallenberg Global Learning Network

Reverberation Transfer Function



- Three sources
- One listener (two ears)
- Filters should include *pinnae filtering* (*spatialized reflections*)
- Filters change if *anything* in the room changes

In principle, this is an exact computational model.

Implementation

Let $h_{ij}(n)$ = impulse response from source j to ear i . Then the output is given by *six convolutions*:

$$y_1(n) = (s_1 * h_{11})(n) + (s_2 * h_{12})(n) + (s_3 * h_{13})(n)$$

$$y_2(n) = (s_1 * h_{21})(n) + (s_2 * h_{22})(n) + (s_3 * h_{23})(n)$$

- For small n , filters $h_{ij}(n)$ are *sparse*
- Tapped Delay Line (TDL) a natural choice

Transfer-function matrix:

$$\begin{bmatrix} Y_1(z) \\ Y_2(z) \end{bmatrix} = \begin{bmatrix} H_{11}(z) & H_{12}(z) & H_{13}(z) \\ H_{21}(z) & H_{22}(z) & H_{23}(z) \end{bmatrix} \begin{bmatrix} S_1(z) \\ S_2(z) \\ S_3(z) \end{bmatrix}$$

Complexity of Exact Reverberation

Reverberation time is typically defined as t_{60} , the time, in seconds, to decay by 60 dB.

Example:

- Let $t_{60} = 1$ second
- $f_s = 50$ kHz
- Each filter h_{ij} requires 50,000 multiplies and additions per sample, or 2.5 *billion* multiply-adds per second.
- Three sources and two listening points (ears) \Rightarrow 30 billion operations per second
 - 10 dedicated CPUs clocked at 3 Gigahertz
 - multiply and addition initiated each clock cycle
 - no wait-states for parallel input, output, and filter coefficient accesses
- FFT convolution is faster, if throughput delay is tolerable

Conclusion: Exact implementation of point-to-point transfer functions is generally too expensive for real-time computation.

Possibility of a Physical Reverb Model

In a complete *physical model* of a room,

- sources and listeners can be moved without affecting the room simulation itself,
- *spatialized* (in 3D) stereo output signals can be extracted using a “virtual dummy head”

How expensive is a room physical model?

- Audio bandwidth = 20 kHz \approx 1/2 inch wavelength
- Spatial samples every 1/4 inch or less
- A 12'x12'x8' room requires $>$ 100 million grid points
- A lossless 3D finite difference model requires one multiply and 6 additions per grid point \Rightarrow 30 billion additions per second at $f_s = 50$ kHz
- A 100'x50'x20' concert hall requires more than *3 quadrillion operations per second*

Conclusion: Fine-grained physical models are too expensive for real-time computation, especially for large halls.

Perceptual Aspects of Reverberation

Artificial reverberation is an unusually interesting signal processing problem:

- “Obvious” methods based on physical modeling or input-output modeling are too expensive
- We do not perceive the full complexity of reverberation
- What is important perceptually?
- How can we simulate only what is audible?

Perception of Echo Density and Mode Density

- For typical rooms
 - Echo density increases as t^2
 - Mode density increases as f^2
- Beyond some time, the echo density is so great that a *stochastic process* results
- Above some frequency, the mode density is so great that a *random frequency response* results
- There is no need to simulate many echoes per sample
- There is no need to implement more resonances than the ear can hear

Proof that Echo Density Grows as Time Squared

Consider a single spherical wave produced from a point source in a rectangular room.

- Tessellate 3D space with copies of the original room
- Count rooms intersected by spherical wavefront

Proof that Mode Density Grows as Freq. Squared

The resonant modes of a rectangular room are given by

$$k^2(l, m, n) = k_x^2(l) + k_y^2(m) + k_z^2(n)$$

- $k_x(l) = l\pi/L_x = l$ th harmonic of the fundamental standing wave in the x
- $L_x =$ length of the room along x
- Similarly for y and z
- Mode frequencies map to a uniform 3D Cartesian grid indexed by (l, m, n)
- Grid spacings are π/L_x , π/L_y , and π/L_z in x, y , and z , respectively.
- Spatial frequency k of mode $(l, m, n) = \text{distance}$ from the $(0,0,0)$ to (l, m, n)
- Therefore, the number of room modes having a given spatial frequency grows as k^2

Early Reflections and Late Reverb

Based on limits of perception, the impulse response of a reverberant room can be divided into two segments

- *Early reflections* = relatively sparse first echoes
- *Late reverberation*—so densely populated with echoes that it is best to characterize the response *statistically*.

Similarly, the *frequency response* of a reverberant room can be divided into two segments.

- Low-frequency sparse distribution of resonant modes
- Modes packed so densely that they merge to form a *random frequency response* with regular statistical properties

Perceptual Metrics for Ideal Reverberation

Some desirable controls for an artificial reverberator include

- $t_{60}(f)$ = desired reverberation time at each frequency
- $G^2(f)$ = signal power gain at each frequency
- $C(f)$ = “clarity” = ratio of impulse-response energy in early reflections to that in the late reverb
- $\rho(f)$ = *inter-aural correlation coefficient* at left and right ears

Perceptual studies indicate that reverberation time $t_{60}(f)$ should be independently adjustable in at least *three* frequency bands.

Energy Decay Curve (EDC)

For measuring and defining reverberation time t_{60} , Schroeder introduced the so-called *energy decay curve (EDC)* which is the *tail integral* of the squared impulse response at time t :

$$\text{EDC}(t) \triangleq \int_t^{\infty} h^2(\tau) d\tau$$

- $\text{EDC}(t)$ = total signal energy remaining in the reverberator impulse response at time t
- EDC decays more smoothly than the impulse response itself
- Better than ordinary amplitude envelopes for estimating t_{60}

Energy Decay Relief (EDR)


The *energy decay relief (EDR)* generalizes the EDC to multiple frequency bands:

$$\text{EDR}(t_n, f_k) \triangleq \sum_{m=n}^M |H(m, k)|^2$$

where $H(m, k)$ denotes bin k of the short-time Fourier transform (STFT) at time-frame m , and M is the number of frames.


- FFT window length $\approx 30 - 40$ ms
- $\text{EDR}(t_n, f_k)$ = total signal energy remaining at time t_n sec in frequency band centered at f_k

Energy Decay Relief (EDR) of a Violin Body Impulse Response



- Energy summed over frequency within each “critical band of hearing” (Bark band)
- Violin body = “small box reverberator”

Reverb = Early Reflections + Late Reverb



- TDL taps may include lowpass filters (air absorption, lossy reflections)
- Several taps may be fed to late reverb unit, especially if it takes a while to reach full density
- Some or all early reflections can usually be worked into the delay lines of the late-reverberation simulation (transposed tapped delay line)

Early Reflections

The “early reflections” portion of the impulse response is defined as everything up to the point at which a statistical description of the late reverb takes hold.


- Often taken to be the first 100ms
- Better to test for *Gaussianness*
 - *Histogram* test for sample amplitudes in 10ms windows
 - *Exponential fit* (t_{60} match) to EDC (Prony’s method, matrix pencil method)
 - *Crest factor* test (peak/rms)
- Typically implemented using *tapped delay lines* (TDL)
(suggested by Schroeder in 1970 and implemented by Moorer in 1979)
- Early reflections should be *spatialized* (Kendall)
- Early reflections influence *spatial impression*

Late Reverberation

Desired Qualities:

1. a smooth (but not too smooth) decay, and
 2. a smooth (but not too regular) frequency response.
- Exponential decay no problem
 - Hard part is making it *smooth*
 - Must not have “flutter,” “beating,” or unnatural irregularities
 - Smooth decay generally results when the echo density is sufficiently high
 - Some short-term energy fluctuation is required for naturalness
 - A smooth *frequency response* has no large “gaps” or “hills”
 - Generally provided when the mode density is sufficiently large
 - Modes should be spread out uniformly
 - Modes may not be too regularly spaced, since audible periodicity in the time-domain can result
 - Moorer’s ideal late reverb: *exponentially decaying white noise*
 - Good smoothness in both time and frequency domains
 - High frequencies need to decay faster than low frequencies
 - Schroeder’s rule of thumb for echo density in the late reverb is 1000 echoes per second or more
 - For impulsive sounds, 10,000 echoes per second or more may be necessary for a smooth response

Schroeder Allpass Sections



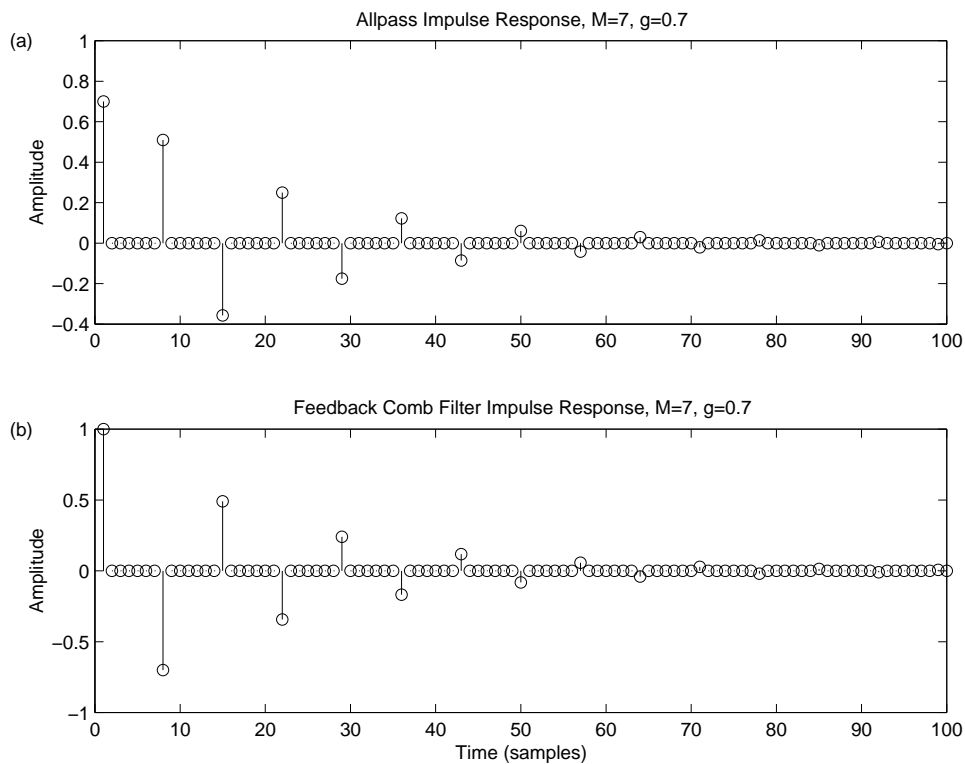
- Typically, $g = 0.7$
- Delay-line lengths M_i mutually prime, and span successive orders of magnitude
e.g., 1051, 337, 113
- Allpass filters in series are allpass
- Each allpass *expands* each nonzero input sample from the previous stage into an entire infinite allpass impulse response
- Allpass sections may be called “*impulse expanders*”, “*impulse diffusers*” or simply “*diffusers*”
- NOT a physical model of diffuse reflection, but single reflections are expanded into many reflections, which is qualitatively what is desired.

Why Allpass?

- Allpass filters do not occur in natural reverberation!
- “Colorless reverberation” is an idealization only possible in the “virtual world”
- **Perceptual factorization:**
Coloration now orthogonal to decay time and echo density

Are Allpasses Really Colorless?

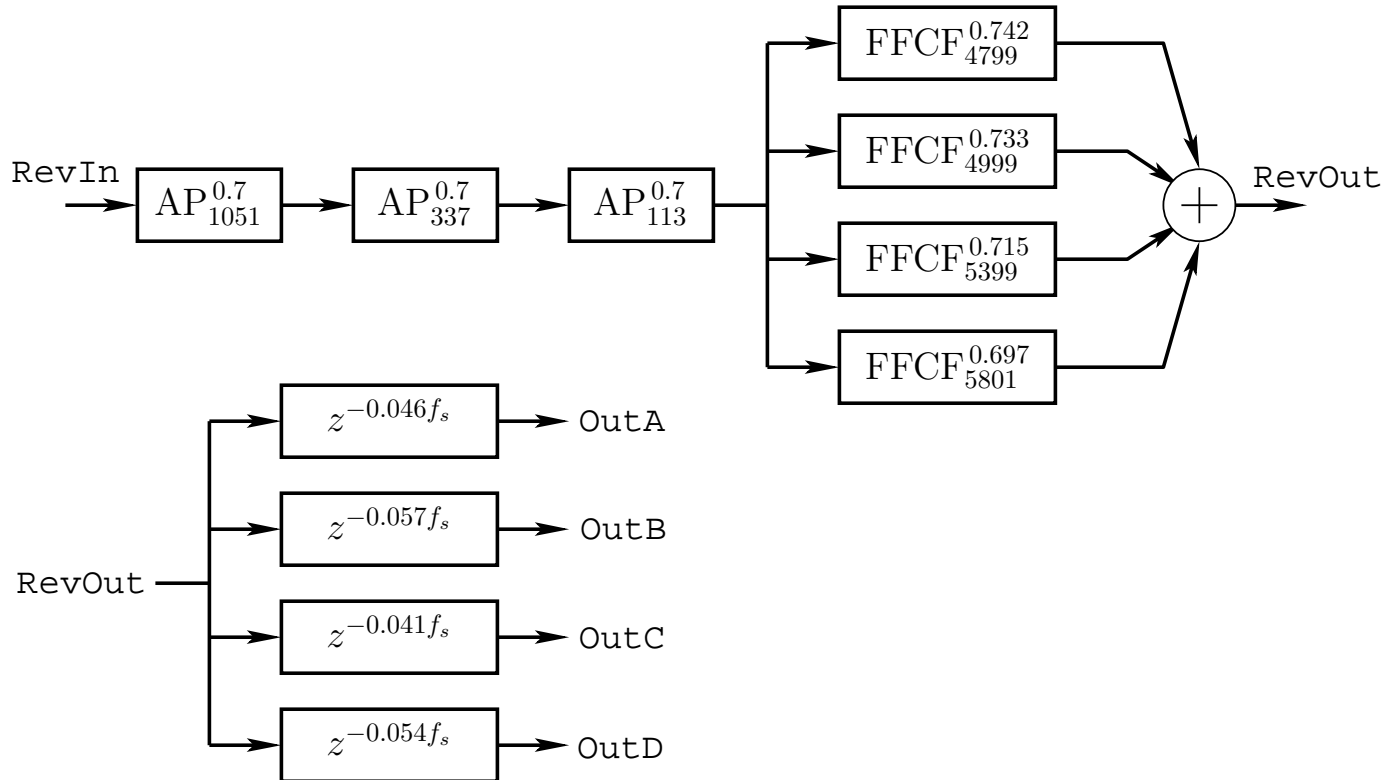
- Allpass impulse response only “colorless” when extremely short (less than 10 ms or so).
- Long allpass impulse responses sound like feedback comb-filters
- The difference between an allpass and feedback-comb-filter impulse response is *one echo!*



$$(a) H(z) = \frac{0.7+z^{-7}}{1+0.7z^{-7}} \quad (b) H(z) = \frac{1}{1+0.7z^{-7}}$$

- Steady-state tones (sinusoids) really do see the same gain at every frequency in an allpass, while a comb filter has widely varying gains.

A Schroeder Reverberator called JCRRev



Classic Schroeder reverberator JCRRev.

JCRRev was developed by John Chowning and others at CCRMA based on the ideas of Schroeder.

- Three Schroeder allpass sections:

$$AP_N^g \triangleq \frac{g + z^{-N}}{1 + gz^{-N}}$$

- Four feedforward comb-filters:

$$FFCF_N^g \triangleq g + z^{-N}$$

- Schroeder suggests a progression of delays close to

$$M_i T \approx \frac{100 \text{ ms}}{3^i}, \quad i = 0, 1, 2, 3, 4.$$

- Comb filters impart distinctive coloration:
 - Early reflections
 - Room size
 - Could be one tapped delay line
- Usage: Instrument adds scaled output to RevIn
- Reverberator output RevOut goes to four *delay lines*
 - Four channels *decorrelated*
 - *Imaging* of reverberation between speakers avoided
- For stereo listening, Schroeder suggests a *mixing matrix* at the reverberator output, replacing the decorrelating delay lines
- A mixing matrix should produce maximally rich yet uncorrelated output signals
- JCRRev is in the Synthesis Tool Kit (STK)
 - JCRRev.cpp
 - JCRRev.h.

FDN Late Reverberation

History

- Gerzon 1971: “orthogonal matrix feedback reverb”
cross-coupled feedback comb filters (see below)
- Stautner and Puckette 1982:

$$\mathbf{A} = g \frac{1}{\sqrt{2}} \begin{bmatrix} 0 & 1 & 1 & 0 \\ -1 & 0 & 0 & -1 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & -1 & 0 \end{bmatrix}$$

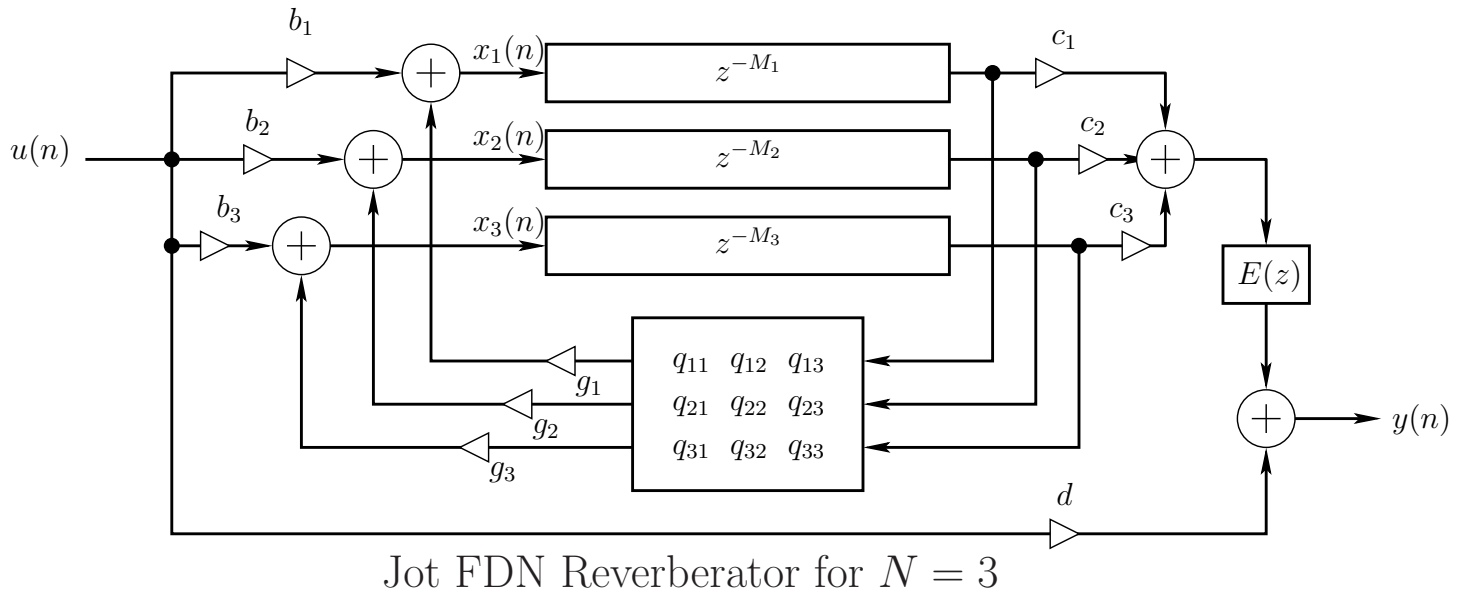
- A second-order *Hadamard matrix*:

$$\mathbf{H}_2 \triangleq \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix},$$

- Higher order Hadamard matrices defined by recursive embedding:

$$\mathbf{H}_4 \triangleq \frac{1}{\sqrt{2}} \begin{bmatrix} \mathbf{H}_2 & \mathbf{H}_2 \\ -\mathbf{H}_2 & \mathbf{H}_2 \end{bmatrix}.$$

Jot's FDN Late Reverberators (1991)



- Generalized state-space model (unit delays replaced by arbitrary delays)
- Note direct path weighted by d
- The “tonal correction” filter $E(z)$ equalizes mode energy independent of reverberation time (perceptual orthogonalization)

Choice of Feedback Matrix

Late reverberation should resemble exponentially decaying noise. This suggests the following two-step procedure for reverberator design:

1. Set $t_{60} = \infty$ and make a *good white-noise generator*
2. Establish desired reverberation times in each frequency band by *introducing losses*

The white-noise generator is the *lossless prototype* reverberator.

Householder Feedback Matrix

Jot proposed the lossless feedback matrix

$$\mathbf{A}_N = \mathbf{I}_N - \frac{2}{N} \underline{u}_N \underline{u}_N^T$$

- *Householder reflection* (negated)
- Input vector is reflected about $\underline{u}_N^T = [1, 1, \dots, 1]$ in N -dimensional space
- \mathbf{I}_N can be replaced by any $N \times N$ permutation matrix
- *Multiply-free* when N is a power of 2
- At most one multiply required
- Only $2N - 1$ additions
(a general matrix-times-vector multiplication is $O(N^2)$ multiply-adds)

Householder Reflection

Let $\mathbf{P}_{\underline{u}}$ denote the *projection matrix* which orthogonally projects vectors onto \underline{u} , i.e.,

$$\mathbf{P}_{\underline{u}} = \frac{\underline{u} \underline{u}^T}{\underline{u}^T \underline{u}} = \frac{\underline{u} \underline{u}^T}{\|\underline{u}\|^2}$$

and

$$\mathbf{P}_{\underline{u}} \underline{x} = \underline{u} \frac{\langle \underline{u}, \underline{x} \rangle}{\|\underline{u}\|^2}$$

specifically projects \underline{x} onto \underline{u} . Since the projection is orthogonal, we have $\langle \underline{x} - \mathbf{P}_{\underline{u}} \underline{x}, \underline{x} \rangle = 0$.

- We may interpret $(I - \mathbf{P}_{\underline{u}}) \underline{x}$ as the *difference vector* between \underline{x} and $\mathbf{P}_{\underline{u}} \underline{x}$, its orthogonal projection onto \underline{u} , since

$$(I - \mathbf{P}_{\underline{u}}) \underline{x} + \mathbf{P}_{\underline{u}} \underline{x} = \underline{x}$$

and we have $(I - \mathbf{P}_{\underline{u}}) \underline{x} \perp \underline{x}$ by definition of the orthogonal projection.


- Consequently, the projection onto \underline{u} *minus* this difference vector gives a *reflection* of the vector \underline{x} about \underline{u} :

$$\underline{y} = \mathbf{P}_{\underline{u}} \underline{x} - (I - \mathbf{P}_{\underline{u}}) \underline{x} = (2\mathbf{P}_{\underline{u}} - I) \underline{x}$$

- \underline{y} is obtained by *reflecting* \underline{x} about \underline{u}
- This is called a *Householder reflection*

Householder FDN = Coupled Feedback Combs

A Householder FDN can be drawn as



- N feedback comb filters in parallel
- Extra global feedback path added, gain = $-2/N$
- Cross-coupled feedback comb filters

Householder Properties for Specific Sizes

- For $N \neq 2$, all entries in the matrix are nonzero
 - Every delay line feeds back to every other delay line
 - Echo density maximized as soon as possible
- For $N = 4$, all matrix entries have the *same magnitude*:

$$\mathbf{A}_4 = \frac{1}{2} \begin{bmatrix} 1 & -1 & -1 & -1 \\ -1 & 1 & -1 & -1 \\ -1 & -1 & 1 & -1 \\ -1 & -1 & -1 & 1 \end{bmatrix}$$

- Only the $N = 4$ case is “balanced” in this way
- Multiply free
- In a manner analogous to Hadamard embedding to generate higher-order Hadamard matrices, Jot proposed constructing an $N = 16$ feedback matrix as a 4×4 *Householder embedding* of the $N = 4$ Householder matrix:

$$\mathbf{A}_{16} = \frac{1}{2} \begin{bmatrix} \mathbf{A}_4 & -\mathbf{A}_4 & -\mathbf{A}_4 & -\mathbf{A}_4 \\ -\mathbf{A}_4 & \mathbf{A}_4 & -\mathbf{A}_4 & -\mathbf{A}_4 \\ -\mathbf{A}_4 & -\mathbf{A}_4 & \mathbf{A}_4 & -\mathbf{A}_4 \\ -\mathbf{A}_4 & -\mathbf{A}_4 & -\mathbf{A}_4 & \mathbf{A}_4 \end{bmatrix}$$

Triangular Feedback Matrices

A triangular matrix has its eigenvalues along the diagonal.

Example:

$$\mathbf{A}_3 = \begin{bmatrix} \lambda_1 & 0 & 0 \\ a & \lambda_2 & 0 \\ b & c & \lambda_3 \end{bmatrix}$$

is *lower triangular*. Its eigenvalues are $(\lambda_1, \lambda_2, \lambda_3)$ for all a, b, c .

Note: Not all triangular matrices with unit-modulus eigenvalues are lossless.

Example:

$$\mathbf{A}_2 = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$$

- Two eigenvalues equal to 1
- Only one eigenvector, $[0, 1]^T$
- Jordan block of order 2 corresponding to the repeated eigenvalue $\lambda = 1$

By direct computation,

$$\mathbf{A}_2^n = \begin{bmatrix} 1 & 0 \\ n & 1 \end{bmatrix}$$

which is clearly not lossless.

Choice of Delay Lengths

- Delay line lengths M_i typically *mutually prime*
- For *sufficiently high mode density*, $\sum_i M_i$ must be sufficiently large.
 - No “ringing tones” in the late impulse response
 - No “flutter”

Mean Free Path

$$\bar{d} = 4\frac{V}{S} \quad (\text{mean free path})$$

where V is the total volume of the room, and S is total surface area enclosing the room.

Regarding each delay line as a mean-free-path delay, the mean free path length, in samples, is the average delay-line length:

$$\frac{\bar{d}}{cT} = \frac{1}{N} \sum_{i=1}^N M_i$$

where c = sound speed and T = sampling period.

This is only a lower bound because many reflections are *diffuse* in real rooms, especially at high frequencies (one plane-wave reflection scatters in many directions)

Mode Density Requirement

FDN order = sum of delay lengths:

$$M \triangleq \sum_{i=1}^N M_i \quad (\text{FDN order})$$

- Order = number of poles
- All M poles are on the unit circle in the lossless prototype
- If uniformly distributed, mode density =

$$\frac{M}{f_s} = MT \quad \text{modes per Hz}$$

- Schroeder suggests that 0.15 modes per Hz (when $t_{60} = 1$ second)
- Generalizing:

$$M \geq 0.15t_{60}f_s$$

- Example: For $f_s = 50$ kHz and $t_{60} = 1$ second, $M \geq 7500$
- Note that $M = t_{60}f_s$ is the length of the FIR filter giving an exact implementation. Thus, recursive filtering is about 7 times more efficient by this rule of thumb.

Achieving Desired Reverberation Times

To set the reverberation time, we need to move the poles of the lossless prototype slightly *inside* the unit circle.

We want the to move high-frequency poles farther in than low-frequency poles.

Basic substitution:

$$z^{-1} \leftarrow G(z)z^{-1}$$

where $G(z)$ a lowpass filter satisfying $|G(e^{j\omega T})| \leq 1$ for all ω .

- $G(z) =$ *per-sample filter* in the propagation medium
First applied to complete reverberators by Jot
- Jot suggests

All pole radii in the reverberator should vary smoothly with frequency.

Otherwise, late decay will be dominated by largest pole(s)

Delay-Filter Design

Let

- $t_{60}(\omega)$ = desired reverberation time at frequency ω
- $H_i(z)$ = lowpass filter for delay-line i

How do we design $H_i(z)$ to achieve $t_{60}(\omega)$?

Let

$$p_i \triangleq e^{j\omega_i T}$$

denote the i th pole of the lossless prototype. Neglecting phase in the loss filter $G(z)$, the substitution

$$z^{-1} \leftarrow G(z)z^{-1}$$

only affects the pole radius, not angle.

Assuming $G(e^{j\omega T}) \approx 1$, pole i moves from $z = e^{j\omega_i T}$ to

$$p_i = R_i e^{j\omega_i T}$$

where

$$R_i = G(R_i e^{j\omega_i T}) \approx G(e^{j\omega_i T}).$$

Desired Pole Radius

Pole radius R_i and t_{60} are related by

$$R_i^{t_{60}(\omega_i)/T} = 0.001$$

The ideal loss filter $G(z)$ therefore satisfies

$$|G(\omega)|^{t_{60}(\omega)/T} = 0.001$$

The desired delay-line filters are therefore

$$H_i(z) = G^{M_i}(z)$$

\Rightarrow

$$|H_i(e^{j\omega T})|^{\frac{t_{60}(\omega)}{M_i T}} = 0.001.$$

or

$$\boxed{20 \log_{10} |H_i(e^{j\omega T})| = -60 \frac{M_i T}{t_{60}(\omega)}}.$$

Now use `invfreqz` or `stmcb`, etc., in Matlab to design low-order filters $H_i(z)$ for each delay line.

First-Order Delay-Filter Design

Jot used first-order loss filters for each delay line:

$$H_i(z) = g_i \frac{1 - a_i}{1 - a_i z^{-1}}$$

- g_i gives desired reverberation time at dc
- a_i sets reverberation time at high frequencies

Design formulas:

$$g_i = 10^{-3M_i T / t_{60}(0)}$$
$$a_i = \frac{\ln(10)}{4} \log_{10}(g_i) \left(1 - \frac{1}{\alpha^2} \right)$$

where

$$\alpha \triangleq \frac{t_{60}(\pi/T)}{t_{60}(0)}$$

Tonal Correction Filter

Let $h_k(n)$ = impulse response of k th system pole. Then

$$\mathcal{E}_k = \sum_{n=0}^{\infty} |h_k(n)|^2 = \text{total energy}$$

Thus, *total energy is proportional to decay time.*

To compensate, Jot proposes a *tonal correction filter* $E(z)$ for the late reverb (not the direct signal).

First-order case:

$$E(z) = \frac{1 - bz^{-1}}{1 - b}$$

where

$$b = \frac{1 - \alpha}{1 + \alpha}$$

and

$$\alpha \triangleq \frac{t_{60}(\pi/T)}{t_{60}(0)}$$

as before.

Further Extensions for FDN Reverberation

While FDNs address several problems in previous reverberation filters, the following areas could benefit from further attention:

- Spatialization of reverberant echoes
 - HRTF for initial early reflection (Kendall and Martens)
 - Diffuse field illusion for late reverb?
(Current approach is simply decorrelating each channel.)
- Mode frequency distribution
 - Coupled delay-line systems tend to have *uniform* mode density
 - Natural mode densities increase with freq. squared
 - However, note that *perception* of mode density *decreases* with frequency
 - Is a uniform distribution a good compromise between nature and perception?

Further Extensions for FDN Reverberation


- Specific low-frequency modes (“early resonances”) — do we want them?
 - Early reflections do some shaping of low-frequency response
 - For concert halls, low-frequency resonances and anti-resonances are typically *eliminated* as much as possible by *parametric equalizers*
 - Perhaps “unnaturally flat” at low frequencies is ideal
 - ‘Small-box reverberators’ such as the violin body or vocal tract require faithful LF resonators (formants important)
- Perceptual studies say reverberation time should be independently adjustable in at least three frequency bands
- Householder FDN = digital waveguide network (DWN) with one scattering junction

FDNs as Digital Waveguide Networks

Householder FDN

$$\mathbf{A}_N = \mathbf{I}_N - \frac{2}{N} \underline{u} \underline{u}^T$$

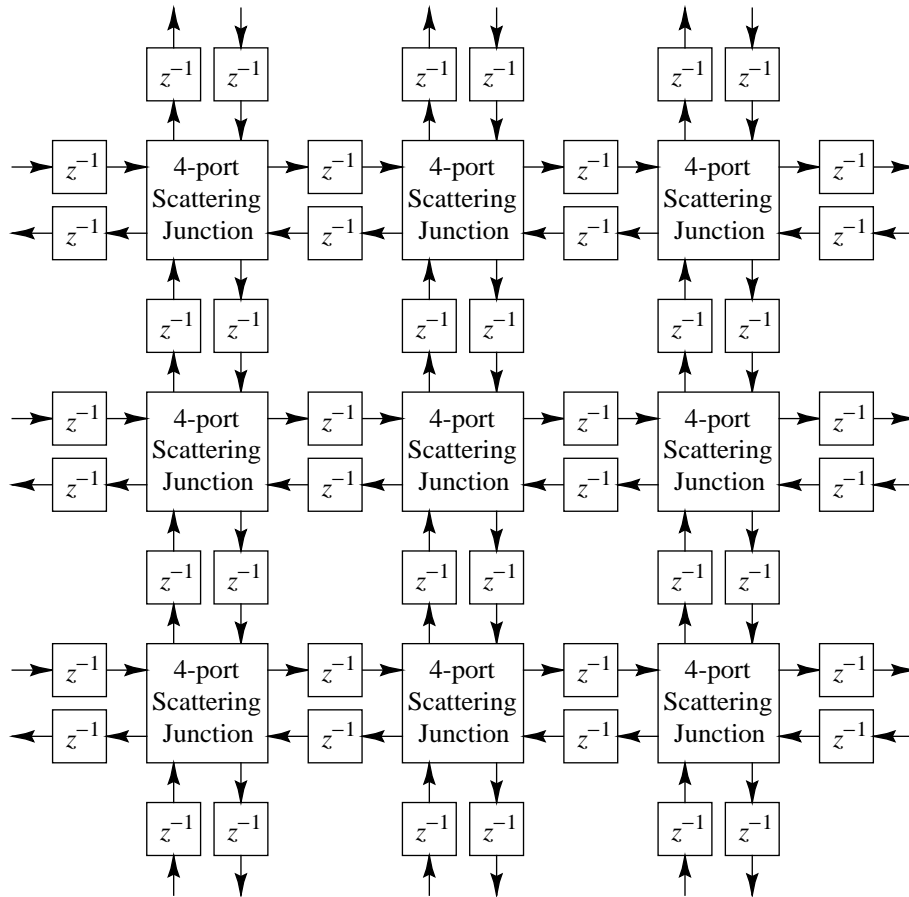

is equivalent to a network of N digital waveguides intersecting at a single scattering junction:



- Single scattering junction indicated by open circle
- Far end of each waveguide branch is terminated by an ideal reflection (filled circle).
- Wave impedance of i th waveguide = $\underline{u}[i]$

- $\underline{u}^T = [1, 1, \dots, 1]$ means all waveguides have same impedance (“isotropic junction”)

Rectilinear Digital Waveguide Mesh



Waveguide Mesh Features

- A *mesh* of such waveguides in 2D or 3D can simulate waves traveling in *any* direction in the space.
- Analogy: tennis racket = rectilinear mesh of strings = pseudo-membrane
- Wavefronts are explicitly simulated in *all directions*
- True *diffuse field* in late reverb
- Spatialized reflections are “free”
- Echo density grows naturally with time
- Mode density grows naturally with frequency
- Low-frequency modes very accurately simulated
- High-frequency modes mistuned due to *dispersion* (can be corrected) (often not heard)
- Multiply free almost everywhere
- Coarse mesh captures most perceptual details

Reverb Resources on the Web

- Harmony Central article¹ (with sound examples)
- William Gardner's MIT Master's thesis²

¹<http://www.harmony-central.com/Effects/Articles/Reverb/>

²<http://www.harmony-central.com/Computer/Programming/virtual-acoustic-room.ps.gz>