Circulant and Elliptic Feedback Delay Networks for Artificial Reverberation

Davide Rocchesso and Julius O. Smith

***Abstract*— The feedback delay network (FDN) has been proposed for digital reverberation. The digital waveguide network (DWN) is also proposed with similar advantages. This paper notes that the commonly used FDN with an $N \times N$ orthogonal feedback matrix is isomorphic to a normalized digital waveguide network consisting of one scattering junction joining N reflectively terminated branches. Generalizations of FDN's and DWN's are discussed. The general case of a lossless FDN feedback matrix is shown to be any matrix having unit-modulus eigenvalues and linearly independent eigenvectors. A special class of FDN's using *circulant* matrices is proposed. These structures can be efficiently implemented and allow control of the time and frequency behavior. Applications of circulant feedback delay networks in audio signal processing are discussed.**

I. Introduction

ARTIFICIAL reverberation is a challenging application in signal processing because it is necessary to approximate large systems (such as concert halls) having hundreds of thousands of poles and zeros in the audio band. Instead of pursuing explicit models that are prohibitively complex, it is necessary to find alternative abstractions that can be implemented at reasonable cost and capture the salient psychoacoustical attributes of natural reverberation. An important practical requirement is a stable numerical implementation of sparse, high-order, nearly lossless linear systems. This paper addresses this and related issues.

A. Prior Work

The field of digital artificial reverberation was launched by Schroeder more than 30 years ago [1]. In his pioneering work, he introduced recursive comb filters and allpass filters as suitable means for inexpensive simulation of multiple echoes. In particular, he introduced use of allpass filters of the form $y(n) = gx(n) + x(n-N) - gy(n-N)$, with N any positive integer, for achieving dense echoes with a flat amplitude response. This structure has since been used extensively in artificial reverberation [2].

In the 1970's, Gerzon [3] generalized the single-input, single-output Schroeder allpass to M inputs and outputs by replacing the N-sample delay line with an order M "unitary network" (a square matrix transfer function having a frequency response matrix that is a unitary matrix at all frequencies, i.e., it must be a "paraunitary" transfer function matrix [4]).

Stautner and Puckette [5] introduced what we now call *feedback delay networks* (FDN's) as structures that are well suited for artificial reverberation. These structures are characterized by a set of delay lines connected in a feedback loop through a "feedback matrix" (see Fig. 1). The FDN was obtained as a generalization of the recursive feedback comb filter $y(n) = x(n-N) + gy(n-N)$ by 1) replacing the single N-sample delay line by a diagonal matrix of delay lines of different lengths and 2) replacing the feedback gain g by the matrix $G = UD$, where U is any unitary matrix,[1] and D is any diagonal matrix having all elements less than $1-\epsilon$ in magnitude, where $\epsilon > 0$ determines the stability margin. Specific early reflections were implemented by adding scaled copies of the source signal into selected points along the delay lines, corresponding to use of the transposed form of the FIR filter [6]. Early reflections in artificial reverberation were apparently first implemented by Moorer using a direct-form FIR filter in series with Schroeder allpass filters and air-absorption comb filters [2].

More recently, Jot has extensively studied FDN's and developed associated techniques for designing good quality reverberators. He suggested the use of efficient special cases of unitary feedback matrices as well as techniques for pole placement to obtain a desired decay-time versus frequency [7] and introduced the valuable design principle that for smoothest (idealized) late reverberation, all modes in a given frequency band should decay at substantially the same rate in order to avoid isolated ringing modes in the late reverberation that tend to sound "metallic" [8].

In 1986, *digital waveguide networks* (DWN) were proposed as a useful starting point for digital reverberator design [9]. The idea was to build an arbitrary closed network of digital waveguides exhibiting the desired early reflections and late echo density and then introduce loss filters into the network to achieve the desired decay time versus frequency. Approaching reverberation via lossless prototypes leads to good numerical and stability properties [4], [10]. Like FDN's, DWN's make it easy to construct well-behaved, high-order, nearly lossless systems.

Manuscript received January 6, 1995; revised July 26, 1996. The associate editor coordinating the review of this paper and approving it for publication was Dr. James H. Snyder.

D. Rocchesso is with the Centro di Sonologia Computazionale, Dipartimento di Elettronica e Informatica, Università degli Studi di Padova, via Gradenigo, 6/A - 35131 Padova, Italy.

J. O. Smith is with the Center for Computer Research in Music and Acoustics (CCRMA), Music Department, Stanford University, Stanford, CA 94305 USA.

Publisher Item Identifier S 1063-6676(97)00758-X.

[1] While papers on this subject speak of unitary feedback matrices $U \in \mathcal{C}^N$ defined by $U^*U = I$, where U^* denotes the Hermitian transpose of U, all practical applications thus far appear to be confined to *orthogonal* matrices $U \in \mathcal{R}^N$.

1063–6676/97$10.00 © 1997 IEEE

B. Lossless Reveberator Prototypes

Reverberator design is often factored into separate designs of the early reflections and the late "diffuse" reverberation. Specific early reflections are easily realized using FIR filter taps. To develop a high-quality late reverberation, it is good to begin with a lossless prototype reverberator so that the structure of its late time response can be clearly heard. Nominally, a lossless prototype reverberator is judged by the quality of *white noise* it generates in response to an impulse signal. For smooth late reverberation, the white noise should sound uniform in every respect. Subsequent introduction of sparse lowpass filters into the prototype network serves to set the desired reverberation time versus frequency. In other words, starting with lossless networks allows the *decoupling* of reverberation time from structural aspects of the reverberator. Any "passive" arithmetic scheme, such as magnitude truncation, can be used at certain multiplier outputs to eliminate the possibility of limit cycles and overflow oscillations [10].

Since FDN's and DWN's appear to present very different approaches for constructing lossless reverberator prototypes, it is natural to ask what connections may exist between them and whether there may be unique advantages of one over the other.

C. Summary and Outline

In Section II, we briefly review the FDN and discuss some of its algebraic properties. In Section III, we explore connections between FDN's and DWN's: It is shown how a single-junction DWN created by the intersection of N waveguides can be interpreted as an order N FDN; conversely, it is shown that any FDN can be interpreted as a DWN, although its scattering junction is not necessarily physical. We derive general conditions for lossless FDN feedback matrices in which the unitary matrix normally used in FDN's is extended to any matrix having unit-modulus eigenvalues and linearly independent eigenvectors. The extension corresponds to a generalization of signal energy by replacing the L_2 norm with an *elliptic* norm induced by any Hermitian, positive-definite matrix [11]. In Section IV, *circulant feedback* matrices are proposed as good choices for FDN's due to their efficiency and versatility in practice. It is straightforward to control the eigenvalues of circulant feedback matrices, and therefore, they can be optimized to yield the best reverberation according to a specified criteria. Finally, in Section V, we present applications in artificial reverberation and use of more general-purpose resonators.

For generality, we treat the complex case, although real numbers are typically used in practice.

II. Feedback Delay Networks

This section reviews FDN's along the lines indicated by Jot [7] and [8] with some modifications.

As depicted in Fig. 1, an FDN is built using N delay lines, each having a length in seconds given by $\tau_i = m_i T$, where $T = 1/F_s$ is the sampling period. The complete FDN is given

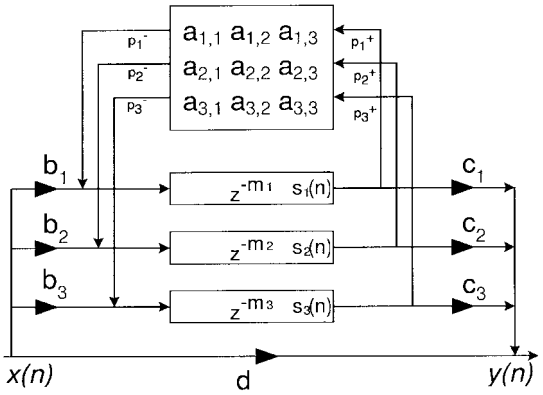

Fig. 1. Order 3 feedback delay network.

by the following relations:

$$y(n) = \sum_{i=1}^{N} c_i s_i(n) + dx(n) \tag{1}$$

$$s_i(n+m_i) = \sum_{j=1}^{N} a_{i,j} s_j(n) + b_i x(n) \tag{2}$$

where $s_i(n)$, $1 \le i \le N$ are the delay-line outputs at time sample n. If $m_i = 1$ for each i, we obtain the conventional state-variable description of a discrete-time linear system [12]. In the present case, m_i are normally large integers; therefore, the variables $s_i(n)$ form only a small subset of the system state at time n, with the remaining state variables being the samples contained within the delay lines at time n. Using the z transform, assuming zero initial conditions, we can rewrite (2) in the frequency domain as

$$Y(z) = \mathbf{c}^T \mathbf{S}(z) + dX(z) \tag{3}$$

$$\mathbf{S}(z) = \mathbf{D}(z)[\mathbf{A}\mathbf{S}(z) + \mathbf{b}X(z)] \tag{4}$$

where $\mathbf{s}^T(z) = [s_1(z), \cdots, s_N(z)]$, $\mathbf{b}^T = [b_1, \cdots, b_N]$, and $\mathbf{c}^T = [c_1, \cdots, c_N]$. The diagonal matrix $\mathbf{D}(z) = \text{diag}\,(z^{-m_1}, z^{-m_2}, \cdots, z^{-m_N})$ is called the "delay matrix," and $\mathbf{A} = [a_{i,j}]_{N \times N}$ is called the "feedback matrix."

The state variables of the FDN can be collected into a vector $\mathbf{w}$ as follows: List the variables contained in the first delay line from the $(m_1 - 1)$th cell to the second cell and then those contained in the second delay line from the $(m_2 - 1)$th cell to the second cell, and so on for the other delay lines; then, attach the output variables s_1 to s_n and, finally, the first cell of all the delay lines in increasing order.

By assuming that each delay line is longer than two samples, the state-variable description corresponding to this variable format for the system (2) can be found to be

$$\begin{aligned} y(n) &= \gamma^T \mathbf{w}(n) + dx(n) \\ \mathbf{w}(n+1) &= \mathbf{A}^\dagger \mathbf{w}(n) + \beta x(n) \end{aligned} \tag{5}$$

where

$$\beta^T = [0, \cdots, 0, \mathbf{b}^T] \tag{6}$$

$$\gamma^T = [0, \cdots, 0, \mathbf{c}^T, \underbrace{0, \cdots, 0}_{N}]. \tag{7}$$

The state-transition matrix is

$$\mathbf{A}^\dagger = \begin{bmatrix} \mathbf{U}_1 & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \mathbf{R}_1 \\ \mathbf{0} & \mathbf{U}_2 & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \mathbf{R}_2 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{U}_N & \mathbf{0} & \mathbf{R}_N \\ \mathbf{P}_1 & \mathbf{P}_2 & \mathbf{P}_3 & \cdots & \mathbf{P}_N & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{A} & \mathbf{0} \end{bmatrix} \tag{8}$$

where

$$\mathbf{U}_j = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & 0 & \cdots & 1 \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix}, \tag{9}$$

$$\mathbf{R}_j = \begin{bmatrix} 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & 0 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ \underbrace{0 \cdots 0}_{j-1} & 1 & 0 & 0 & \cdots & 0 \end{bmatrix} \tag{10}$$

and

$$\mathbf{P}_j = \begin{bmatrix} & & & & & \\ j-1 & \left\{ \begin{matrix} 0 & 0 & 0 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & 0 & \cdots & 0 \end{matrix} \right. & & & & \\ & \begin{matrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & 0 & \cdots & 0 \end{matrix} & & & & \end{bmatrix}. \tag{11}$$

We have

$$\begin{aligned} &\mathbf{U}_j \in \mathcal{C}^{(m_j-2)\times(m_j-2)} \\ &\mathbf{R}_j \in \mathcal{C}^{(m_j-2)\times N} \\ &\mathbf{P}_j \in \mathcal{C}^{N\times(m_j-2)}. \end{aligned} \tag{12}$$

The order of the system (5) is equal to the sum of the delay-line lengths:

$$N^\dagger = \sum_{i=1}^{N} m_i.$$

From (4), the transfer function is easily found to be

$$\begin{aligned} H(z) &= \frac{Y(z)}{X(z)} \\ &= \mathbf{c}^T[\mathbf{D}(z^{-1}) - \mathbf{A}]^{-1}\mathbf{b} + d. \end{aligned} \tag{13}$$

Note that when $m_i = 1$ for all i, the FDN specializes to a fully general state-space description [12]. This implies that *any* linear, time-invariant, discrete-time system can be formulated as a special case of an FDN since every state-space description is a special case. This suggests that a wide variety of stable FDN's can be generated by starting with any stable LTI system whatsoever and performing the substitution $z^{-1} \leftarrow z^{-m_i}$ on each delay element or any other conformal mapping that takes the unit circle to itself (another example being the Schroeder allpass transformation $z^{-1} \leftarrow (\rho + z^{-m_i})/(1 + \rho z^{-m_i})$). Stability is preserved even when the unit-sample delays of the original state-space description are mapped using different conformal maps. This can be seen from the matrix power series expansion

$$[\mathbf{I} - \mathbf{D}(z)\mathbf{A}]^{-1} = \mathbf{I} + \mathbf{D}(z)\mathbf{A} + \cdots + \mathbf{D}^k(z)\mathbf{A}^k + \cdots \tag{14}$$

As long as $||\mathbf{D}(e^{j\omega})|| \leq \mathbf{I}$, by making use of the triangle inequality and the Cauchy–Schwarz (or mutual consistency) inequality [11], we can write

$$\begin{aligned} &||\mathbf{I} + \cdots + \mathbf{D}^k(z)\mathbf{A}^k + \cdots|| \\ &\quad \leq ||\mathbf{I}|| + \cdots + ||\mathbf{D}^k(z)\mathbf{A}^k|| + \cdots \\ &\quad \leq 1 + \cdots + ||\mathbf{D}^k(z)|| \cdot ||\mathbf{A}^k|| + \cdots \\ &\quad \leq 1 + \cdots + ||\mathbf{D}(z)||^k \cdot ||\mathbf{A}||^k + \cdots \\ &\quad \leq 1 + \cdots + ||\mathbf{A}||^k + \cdots \\ &\quad = \frac{1}{1 - ||\mathbf{A}||}. \end{aligned} \tag{15}$$

Therefore, as long as $||A||^n$ decays exponentially with n, stability is assured. The above derivation extends immediately to time-varying feedback matrices $\mathbf{A}_k$, provided $||\mathbf{A}_k|| \leq 1 - \epsilon$ for some worst-case $\epsilon > 0$.

The poles of the FDN are the solutions of either

$$\det[\mathbf{A} - \mathbf{D}(z^{-1})] = 0 \tag{16}$$

or

$$\det[z\mathbf{I} - \mathbf{A}^\dagger] = 0, \qquad z \neq 0. \tag{17}$$

The matrix $\mathbf{A}^\dagger$ is not uniquely determined by $\mathbf{A}$. In fact, our ordering of the state variables differs from that used by Jot [7]. Our ordering gives

$$\mathbf{A}^\dagger \mathbf{A}^{\dagger T} = \text{diag}\,(\mathbf{I}_{m_1-1}, \cdots, \mathbf{I}_{m_N-1}, \mathbf{A}\mathbf{A}^T). \tag{18}$$

From (18), we see that $\mathbf{A}$ is unitary if and only if $\mathbf{A}^\dagger$ is unitary. Since a unitary matrix has eigenvalues on the unit circle, we see from (17) that it is sufficient to choose a unitary matrix in order to have all the system poles on the unit circle. This yields a lossless FDN prototype.

By application of the matrix inversion lemma [12] to the transfer function (13), the system zeros are found as the solutions of

$$\det\left[\mathbf{A} - \mathbf{b}\,\frac{1}{d}\,\mathbf{c}^T - \mathbf{D}(z^{-1})\right] = 0. \tag{19}$$

The formulation of (2) represents a prototype structure in the sense that with the appropriate choice of feedback matrix, it is a lossless structure. In practice, we must insert attenuation coefficients and filters in the feedback loop. For example, one may insert a gain [8]

$$g_i = \alpha^{m_i} \tag{20}$$

at the output of each delay line in the FDN. This corresponds to replacing $D(z)$ with $D(z/\alpha)$ in (4). With this choice of the attenuation coefficients, all the system poles are uniformly contracted by a factor α, thus ensuring a uniform decay of all the modes.

In a practical realization, we normally need to introduce frequency-dependent losses such that higher frequencies decay faster. We can do this by introducing lowpass filters after each

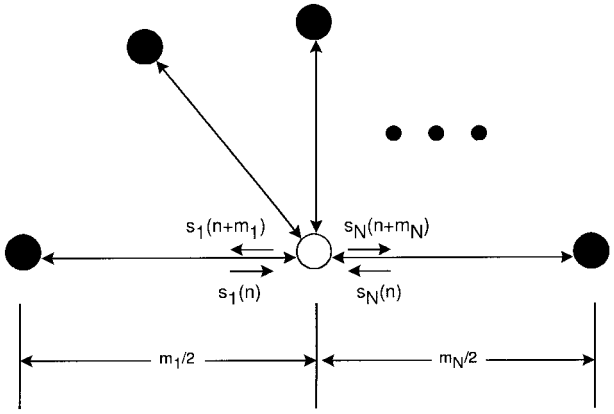

Fig. 2. Waveguide network consisting of a single scattering junction, indicated by an open circle, to which N branches are connected. The far end of each branch is terminated by an ideal, noninverting reflection.

delay line in place of the gains g_i. In this case, local uniformity of mode decay is still achieved by condition (20), where g_i and α are made frequency dependent:

$$G_i(z) = A^{m_i}(z) \tag{21}$$

where $A(z)$ can be interpreted as the per-sample filtering [13], [14], [8].

Notice that uniform decay of all the modes, albeit arguably desirable in artificial reverberators for a smooth late time response, is not found in actual rooms. Normal modes are associated with standing waves, which have an absorption that depends on their orientation. For example, in a rectangular enclosure, waves traveling in a direction normal to a wall are less absorbed than oblique waves [15, p. 392] so that the corresponding standing waves (which are expressible as the superposition of traveling waves in opposite directions) have different reverberation times. The room-acoustics interpretation of FDN's provided in Section V points to ways of modeling such uneven decays.

III. Digital Waveguide Networks

Digital waveguide networks provide a useful paradigm for sound synthesis based on physical modeling [14]. They have also been proposed for constructing arbitrarily complex digital reverberators [9], which are free of limit cycles and overflow oscillations if passive arithmetic is used [10]. In this section, we explore the relationships between DWN's and FDN's.

Fig. 2 illustrates an N-branch DWN that is structurally equivalent to the feedback loop of an Nth-order FDN. It consists of a single scattering junction, indicated by a white circle, to which N branches are connected. The far end of each branch is terminated by an ideal noninverting reflection (black circle). The waves traveling into the junction are associated with the FDN delay line outputs $s_i(n)$, and the length of each waveguide is *half* the length of the corresponding FDN delay line m_i (since a traveling wave must traverse the branch twice to complete a round trip from the junction to the termination and back). When m_i is odd, we may replace the reflecting termination by a unit-sample delay.

A. Lossless Scattering

The delay-line inputs (outgoing traveling waves) are computed by multiplying the delay-line outputs (incoming traveling waves) by the $N \times N$ feedback matrix (scattering matrix) $\mathbf{A} = [a_{i,j}]$. By defining $p_i^+ = s_i(n)$, $p_i^- = s_i(n + m_i)$, we obtain the more usual DWN notation

$$\mathbf{p}^- = \mathbf{A}\mathbf{p}^+ \tag{22}$$

where

- $\mathbf{p}^+$ vector of incoming traveling-wave samples arriving at the junction at time n,
- $\mathbf{p}^-$ vector of outgoing traveling-wave samples leaving the junction at time n,
- $\mathbf{A}$ *scattering matrix* associated with the waveguide junction.

The junction of N physical waveguides determines the structure of the matrix $\mathbf{A}$ according to the basic principles of physics.

Considering the parallel junction of N lossless acoustic tubes, each having characteristic admittance Γ_j, the continuity of pressure and conservation of volume velocity at the junction give us the following scattering matrix for the pressure waves:

$$\mathbf{A} = \begin{bmatrix} \frac{2\Gamma_1}{\Gamma_J} - 1 & \frac{2\Gamma_2}{\Gamma_J} & \cdots & \frac{2\Gamma_N}{\Gamma_J} \\ \frac{2\Gamma_1}{\Gamma_J} & \frac{2\Gamma_2}{\Gamma_J} - 1 & \cdots & \frac{2\Gamma_N}{\Gamma_J} \\ \cdots & & & \\ \frac{2\Gamma_1}{\Gamma_J} & \frac{2\Gamma_2}{\Gamma_J} & \cdots & \frac{2\Gamma_N}{\Gamma_J} - 1 \end{bmatrix} \tag{23}$$

where

$$\Gamma_J = \sum_{i=1}^{N} \Gamma_i. \tag{24}$$

Equation (23) can be derived by first writing the volume velocity at the jth tube in terms of pressure waves as $v_j = (p_j^+ - p_j^-)\Gamma_j$. Applying the conservation of velocity, we can find the expression $p = 2\sum_{i=1}^{N} \Gamma_i p_i^+ / \Gamma_J$ for the junction pressure. Finally, if we express the junction pressure as the sum of incoming and outgoing pressure waves at any branch, we derive (23).

B. Normalized Scattering

For ideal numerical scaling in the L_2 sense, we may choose to propagate *normalized waves* that lead to normalized scattering junctions analogous to those encountered in normalized ladder filters [17]. Normalized waves may be either normalized pressure $\tilde{p}_j^+ = p_j^+ \sqrt{\Gamma_i}$ or normalized velocity $\tilde{v}_j^+ = v_j^+ / \sqrt{\Gamma_i}$. Since the signal power associated with a traveling wave is simply $\mathcal{P}_j^+ = (\tilde{p}_j^+)^2 = (\tilde{v}_j^+)^2$, they may also be called *root-power waves* [10].

The scattering matrix for normalized pressure waves is given by

$$\tilde{\mathbf{A}} = \begin{bmatrix} \frac{2\Gamma_1}{\Gamma_J} - 1 & \frac{2\sqrt{\Gamma_1\Gamma_2}}{\Gamma_J} & \cdots & \frac{2\sqrt{\Gamma_1\Gamma_n}}{\Gamma_J} \\ \frac{2\sqrt{\Gamma_2\Gamma_1}}{\Gamma_J} & \frac{2\Gamma_2}{\Gamma_J} - 1 & \cdots & \frac{2\sqrt{\Gamma_2\Gamma_n}}{\Gamma_J} \\ \cdots & & \cdots & \\ \cdots & & \cdots & \\ \frac{2\sqrt{\Gamma_n\Gamma_1}}{\Gamma_J} & \frac{2\sqrt{\Gamma_n\Gamma_2}}{\Gamma_J} & \cdots & \frac{2\Gamma_n}{\Gamma_J} - 1 \end{bmatrix}. \quad (25)$$

The normalized scattering matrix can be expressed as a Householder reflection

$$\tilde{\mathbf{A}} = \frac{2}{||\tilde{\mathbf{\Gamma}}||^2}\tilde{\mathbf{\Gamma}}\tilde{\mathbf{\Gamma}}^T - \mathbf{I} \quad (26)$$

where $\tilde{\mathbf{\Gamma}}^T = [\sqrt{\Gamma_1}, \cdots, \sqrt{\Gamma_N}]$, and Γ_i is the wave admittance in the ith waveguide branch. The geometric interpretation of (26) is that the incoming pressure waves are reflected about the vector $\tilde{\mathbf{\Gamma}}$. Unnormalized scattering junctions can be expressed in the form of an "oblique" Householder reflection $\mathbf{A} = 2\mathbf{1}\mathbf{\Gamma}^T/\langle \mathbf{1}, \mathbf{\Gamma}\rangle - \mathbf{I}$, where $\mathbf{1}^T = [1, \cdots, 1]$ and $\mathbf{\Gamma}^T = [\Gamma_1, \cdots, \Gamma_N]$.

C. Complexity

It is important to note that a Householder reflection can be implemented using $O(N)$ numerical operations, as opposed to $O(N^2)$ operations for a general scattering matrix (in computing $\mathbf{A}\tilde{p}^+$ in (26), first precompute the inner product $\tilde{\mathbf{\Gamma}}^T\tilde{p}^+$ [18]). Since all junctions of N physical waveguides can be expressed as a Householder reflection, all such scattering junctions require only $O(N)$ computations.

It is interesting to note that Jot [7] proposed a class of feedback matrices for the efficient implementation of FDN's that are specialized Householder reflections. We have just shown that the same kind of structure arises naturally, in the context of waveguide modeling, for physically based scattering matrices.

D. Conditions for Losslessness

The scattering matrices for lossless physical waveguide junctions give an apparently unexplored class of lossless FDN prototypes. However, this is just a subset of all possible lossless feedback matrices. We are therefore interested in the most general conditions for losslessness of an FDN feedback matrix.

Consider the general case in which $\mathbf{A}$ is allowed to be any scattering matrix, i.e., it is associated with a not-necessarily-physical junction of N physical waveguides. Following the definition of losslessness in classical network theory, we may say that a waveguide scattering matrix $\mathbf{A}$ is said to be *lossless* if the *total complex power* [19] at the junction is scattering invariant, i.e.,

$$\begin{aligned} &\mathbf{p}^{+*}\mathbf{\Gamma}\mathbf{p}^+ = \mathbf{p}^{-*}\mathbf{\Gamma}\mathbf{p}^- \\ &\Rightarrow \mathbf{A}^*\mathbf{\Gamma}\mathbf{A} = \mathbf{\Gamma} \end{aligned} \quad (27)$$

where $\mathbf{\Gamma}$ is any Hermitian, positive-definite *matrix* (which has an interpretation as a generalized junction admittance). The form $x^*\mathbf{\Gamma}x$ is by definition the square of the *elliptic norm* of x induced by $\mathbf{\Gamma}$, or $||x||^2_{\mathbf{\Gamma}} = x^*\mathbf{\Gamma}x$. Setting $\mathbf{\Gamma} = \mathbf{I}$, we obtain that $\mathbf{A}$ must be unitary. This is the case commonly used in current FDN practice.

The following theorem gives a general characterization of lossless scattering:

Theorem 1: A scattering matrix (FDN feedback matrix) $\mathbf{A}$ is lossless if and only if its eigenvalues lie on the unit circle and it admits a basis of linearly independent eigenvectors.

Proof: In general, the Cholesky factorization $\mathbf{\Gamma} = \mathbf{U}^*\mathbf{U}$ gives an upper triangular matrix $\mathbf{U}$ that converts $\mathbf{A}$ to a unitary matrix via similarity transformation: $\mathbf{A}^*\mathbf{\Gamma}\mathbf{A} = \mathbf{\Gamma} \Rightarrow \mathbf{A}^*\mathbf{U}^*\mathbf{U}\mathbf{A} = \mathbf{U}^*\mathbf{U} \Rightarrow \tilde{\mathbf{A}}^*\tilde{\mathbf{A}} = \mathbf{I}$, where $\tilde{\mathbf{A}} = \mathbf{U}\mathbf{A}\mathbf{U}^{-1}$. Hence, the eigenvalues of every lossless scattering matrix lie on the unit circle. It readily follows from similarity to $\tilde{\mathbf{A}}$ that $\mathbf{A}$ admits N linearly independent eigenvectors. In fact, $\tilde{\mathbf{A}}$ is a normal matrix (since it is unitary), and normal matrices admit a basis of linearly independent eigenvectors [20].

Conversely, assume that $|\lambda| = 1$ for each eigenvalue of $\mathbf{A}$ and that there exists a matrix $\mathbf{T}$ of linearly independent eigenvectors of $\mathbf{A}$. Then, the matrix $\mathbf{T}$ diagonalizes $\mathbf{A}$ to give $\mathbf{T}^{-1}\mathbf{A}\mathbf{T} = \mathbf{D} \Rightarrow \mathbf{T}^*\mathbf{A}^*\mathbf{T}^{-1*} = \mathbf{D}^*$, where $\mathbf{D} = \text{diag}\,(\lambda_1, \cdots, \lambda_N)$. Multiplying, we obtain $\mathbf{T}^*\mathbf{A}^*\mathbf{T}^{-1*}\mathbf{T}^{-1}\mathbf{A}\mathbf{T} = \mathbf{D}^*\mathbf{D} = \mathbf{I} \Rightarrow \mathbf{A}^*\mathbf{T}^{-1*}\mathbf{T}^{-1}\mathbf{A} = \mathbf{T}^{-1*}\mathbf{T}^{-1}$. Thus, (27) is satisfied for $\mathbf{\Gamma} = \mathbf{T}^{-1*}\mathbf{T}^{-1}$, which is Hermitian and positive definite. □

Thus, lossless scattering matrices may be fully parameterized as $\mathbf{A} = \mathbf{T}^{-1}\mathbf{D}\mathbf{T}$, where $\mathbf{D}$ is any unit-modulus diagonal matrix, and $\mathbf{T}$ is any invertible matrix. In the real case, we have $\mathbf{D} = \text{diag}\,(\pm 1)$ and $\mathbf{T} \in \Re^{N\times N}$.

It can be quickly verified[2] that all scattering matrices arising from the parallel intersection of N physical waveguides possess one eigenvalue equal to 1 and $N-1$ eigenvalues equal to -1 [21]. In the case of physical waveguides of equal impedances, the eigenvector associated with the eigenvalue 1 corresponds to equal incoming waves, whereas an eigenvector associated with the eigenvalue -1 corresponds to equal incoming waves on $N-1$ branches and a large opposite wave on the remaining branch that pulls the junction pressure to zero.

Theorem 1 characterizes lossless FDN feedback matrices $\mathbf{A}$ as those having eigenvalues on the unit circle, where the definition of losslessness was given by (27). It remains to be shown that $\mathbf{A}$ satisfying (27) implies that the poles of the corresponding FDN are all on the unit circle. To this end, recall the form of the state-transition matrix (8), and define the extended generalized admittance

$$\Gamma^\dagger = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Gamma} \end{bmatrix}. \quad (28)$$

By analogy to the derivation of (18), we get

$$\mathbf{A}^{\dagger^T}\Gamma^\dagger\mathbf{A}^\dagger = \text{diag}\,(\mathbf{I}_{m_1-1}, \cdots, \mathbf{I}_{m_N-1}, \mathbf{A}^T\Gamma\mathbf{A}). \quad (29)$$

[2] Using the eigenvectors $\mathbf{e}_0^T = [1, \cdots, 1]$ and $\mathbf{e}_k^T = [1, \cdots, 1, \underbrace{1 - \Gamma_J/\Gamma_k}_{k\text{th}}, 1, \cdots, 1]$, $k = 1, \cdots, N$

This equation shows that $\mathbf{A}$ is lossless if and only if $\mathbf{A}^\dagger$ is lossless, and its eigenvalues are on the unit circle by Theorem 1. However, from (17), we have that the eigenvalues of $\mathbf{A}^\dagger$ are the poles of the corresponding FDN. Therefore, the FDN is lossless, and its impulse response consists only of nonincreasing and nondecaying modes. Desired decay characteristics versus frequency for obtaining a specific artificial reverberator can now be controlled separately by means of attenuation coefficients and filters, as indicated in Section II.

E. Relation of DWN's to FDN's

When $\mathbf{U}$ in the proof of Theorem 1 is diagonal and positive, a physical waveguide interpretation always exists with $\mathbf{U} = \text{diag}\,(\tilde{\mathbf{\Gamma}})$. A generalized waveguide interpretation exists for all $\mathbf{U}$ via vector transformers [16, p. 55, sec. 4] in which $\mathbf{U}$ acts as an ideal transformer (in the classical network theory sense) on the vector of all N waveguide variables. If $\mathbf{p} = \mathbf{p}^+ + \mathbf{p}^-$ denotes the vector of physical pressures at the junction and $\mathbf{v} = \mathbf{v}^+ + \mathbf{v}^-$ denotes the physical volume velocities, then we have that the junction power, which is defined as $\mathcal{P}_J = \mathbf{v}^*\mathbf{p}$, is invariant with respect to insertion of a vector transformer (similarity transformation applied to the scattering matrix).

A unitary FDN feedback matrix corresponds to a not-necessarily-physical scattering junction in which the total complex power is given by the ordinary L_2 norm of the incoming or outgoing traveling waves. Since the physical power associated with an incoming wave vector $\mathbf{p}^+$ is $\mathcal{P}_J = \mathbf{p}^{+*}\mathbf{\Gamma}\mathbf{p}^+$, where $\mathbf{\Gamma} = \text{diag}\,(\Gamma_1, \cdots, \Gamma_N)$, we see that $\mathbf{A}$ unitary corresponds to a scattering junction joining waveguides of equal wave impedance, i.e., $\mathbf{\Gamma} = \text{diag}\,(1, \cdots, 1)$. In the more general (unnormalized) case in which the branch impedances are different, i.e., $\mathbf{\Gamma} = \text{diag}\,(\Gamma_1, \cdots, \Gamma_N)$, we obtain the larger class of scattering matrices that preserve an elliptic norm as induced by a positive-definite (or Hermitian) generalized junction impedance.

Since, as discussed above, only a subset of all $N \times N$ unitary matrices is given by a physical junction of N waveguides, the unitary FDN point of view yields lossless systems outside the scope of those suggested by multiport scattering theory. On the other hand, since only normalized waveguide junctions exhibit unitary scattering matrices, the DWN approach gives rise to new classes of FDN's. Moreover, by considering more than one scattering junction, the DWN approach suggests new classes of network topologies following physical analogies. Similarly, FDN matrices can be partitioned to embed several FDN subsystems into larger FDN systems.

Formally, every DWN can be expressed as an FDN by collecting all of its delay lines into a diagonal delay matrix $\mathbf{D}(z)$ as in (4) and finding the matrix $\mathbf{A}$ that computes the delay-line inputs from the delay-line outputs. Therefore, every waveguide network yields a feedback matrix for consideration in the FDN framework. Conversely, every real lossless FDN can be expressed as a single-junction not-necessarily-physical normalized waveguide network using an ideal vector transformer at the junction of equal-impedance waveguide branches.

F. Finite-Wordlength Effects

We have just seen how the FDN can be seen as a simple DWN having a not-necessarily-physical scattering matrix. In order to provide an easy control over the decay, the scattering matrix has to satisfy the condition of losslessness (27). A finite-precision implementation of the FDN might incur in limit cycles or overflow oscillations due to departures from the infinite-precision lossless prototype. Departures can be of two kinds: The finite-precision scattering matrix does not satisfy the lossless condition (27) or the round-off noise in the matrix by vector multiplication $\mathbf{A}\mathbf{p}^+$ introduces signal amplitude modifications. By assuming that the scattering matrix satisfies (27) in a large extent even in finite precision, it is possible to apply the arguments used in [9], [10], [16], and [22] for the DWN's in order to avoid limit cycles or overflow oscillations. If the matrix by vector multiplication is performed in the straightforward way as a collection of inner products, and the matrix coefficients have the same n bits of precision as the signals, it is sufficient to perform these order-N inner products in the extended precision of $2n+N-1$ bits and apply a passive truncation scheme on the output signal. In two's complement arithmetic, a simple passive truncation scheme is the following:

- If the $N-1$ most significant bits are not equal, replace the output value by the maximum-magnitude number in n-bit two's complement having the correct sign (saturation).
- Discard the n least significant bits, and add 2^{-n+1} to the result if it is negative.

As far as the condition on the losslessness of the scattering matrix is concerned, general requirements for the construction of "structurally lossless," or at least "structurally passive," scattering matrices have to be worked out. This topic, which was previously touched on by Gray [22] in the $N = 2$ case, will be discussed in a forthcoming paper since a complete treatment would enlarge the scope of this paper significantly.

IV. Circulant Feedback Delay Networks

Consider the class of circulant feedback matrices having the form

$$\mathbf{A} = \begin{bmatrix} a(0) & a(1) & \cdots & a(N-1) \\ a(N-1) & a(0) & \cdots & a(N-2) \\ \cdots & & & \\ a(1) & \cdots & a(N-1) & a(0) \end{bmatrix}.$$

This class of matrices gives rise to a class of FDN's we call *circulant feedback delay networks* (CFDN's). The following two facts can be proved [23]:

Fact 1: If a matrix is circulant, it is normal, i.e., $\mathbf{A}^*\mathbf{A} = \mathbf{A}\mathbf{A}^*$.

Fact 2: If a matrix is circulant and lossless, it is unitary.

It is well known that every circulant matrix is diagonalized by the discrete Fourier transform (DFT) matrix [23]. This implies that the eigenvalues of $\mathbf{A}$ can be computed by means of the DFT of the first row:

$$\begin{aligned} \{\lambda(\mathbf{A})\} &= \{\mathbf{A}(k)\} \\ &= \text{DFT}\,\{[a(0) \cdots a(N-1)]^T\} \end{aligned}$$

where $\{\lambda(\mathbf{A})\}$ denotes the set of all eigenvalues of $\mathbf{A}$, and

$\{\mathbf{A}(k)\}$ denotes the set of complex DFT samples obtained from taking the DFT of $\{a(\cdot)\}$.

A. Design of Poles and Zeros in CFDN's

A matrix that is both unitary and circulant has all eigenvalues on the unit circle, and the DFT can be used to compute the eigenvalue phases. In the case of equal-length delay lines, the eigenvalues determine the resonance frequencies in a simple way. From (16), when $\mathbf{D}(z^{-1}) = z^m\mathbf{I}$, the system poles are the mth complex roots of the eigenvalues of $\mathbf{A}$.

Conversely, we can easily design a circulant matrix to have a desired distribution of eigenvalues. This is also true for any lossless matrix since Theorem 1 gives that any $\mathbf{A}$ of the form $\mathbf{A} = \mathbf{T}^{-1}\mathbf{D}\mathbf{T}$ is lossless, where $\mathbf{D}$ is any unit-modulus diagonal matrix, and $\mathbf{T}$ is any invertible matrix. Thus, a lossless matrix is characterized by the arguments of its eigenvalues and a similarity transformation matrix $\mathbf{T}$. The advantage of choosing circulant FDN's over other kinds of FDN's is the possibility of computing $\mathbf{A}$ from its eigenvalues very efficiently by means of a single inverse FFT.

As we will see in Section V, in a practical implementation, the delay lengths are typically not equal. However, the equal-delay case is easier to analyze. The limitations and advantages of such a choice will become clearer in Section V.

The actual presence of resonance peaks corresponding to the eigenvalues depends on the positions of the zeros, as given by (19). Assuming equal-length delay lines and $d = 1$, (19) becomes

$$\det\left[\mathbf{A} - \mathbf{b}\mathbf{c}^T - z^m\mathbf{I}\right] = 0 \tag{30}$$

which means that the zeros are the mth complex roots of the eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{c}^T$.

In order to have "colorless" reverberation, it may be desirable to make the envelope of the amplitude response flat. To do this, each zero should be equal to the reciprocal of a pole. In prototype CFDN's, the feedback matrix is lossless, and the system poles are on the unit circle so that the zeros must equal the poles. However, when all zeros and poles cancel exactly, the impulse response of the FDN degenerates to an impulse.[3] This is a general problem with any allpass reverberator: Lengthening the reverberation time without changing the delay lengths forces the impulse response to converge to an impulse. In our case, we depart from the idealized case by slightly changing the delay lengths. As we will show in Section V, this approach leads to reverberators having a frequency response that is nearly flat at low frequencies while preserving the richness of the echo density in the time domain.

Therefore, we continue treating the prototype case of equal-length delay lines and $d = 1$ and show that we can obtain perfect canceling of zeros and poles by using 1) $\mathbf{b}^T = [1, 1, \cdots, 1]$ and 2) $\mathbf{c}$ having n entries equal to 1, n entries equal to -1, and zeros for the remaining entries. This result is due to the following

[3] Since we are discussing discrete-time systems, the term "impulse" means the same thing as "unit sample pulse."

Theorem 2: Given a circulant $N \times N$ matrix $\mathbf{A}$, let $\mathbf{A}'$ be obtained by adding a constant c to each entry of n rows (columns) and subtracting the same constant c from each entry of another n rows (columns). Then, $\mathbf{A}$ and $\mathbf{A}'$ have the same eigenvalues.

Before providing the proof of Theorem 2, we need to prove the following

Lemma: All the eigenvectors of a circulant matrix other than the "dc" vector $[1, \cdots, 1]^T$ lie in the null space of any matrix with constant rows. Thus, adding constant rows cannot alter eigenvalues or eigenvectors other than the zeroth.

Proof: This follows immediately from the fact that the eigenvectors of every circulant matrix are given by the columns of the DFT matrix of the same size, and these vectors are orthogonal. Therefore, a constant row is orthogonal to all eigenvectors of the DFT matrix except the dc eigenvector.

The lemma states that all we can do by adding constant rows to a circulant matrix is move the "dc" eigenvector to some other vector and change its eigenvalue.

Proof of Theorem 2: Consider the matrix $\mathbf{A}'$ given by

$$\mathbf{A}' = \mathbf{A} - \mathbf{b}\mathbf{c}^T \tag{31}$$

where $\mathbf{b}^T = [1, -1, 0, \cdots, 0]$ and $\mathbf{c}^T = [1, 1, \cdots, 1]$. Since any circulant matrix is diagonalized by the DFT matrix, if we premultiply and postmultiply both sides of (31) by the DFT matrices $\mathbf{F}^*$ and $\mathbf{F}$, we obtain

$$\begin{aligned}\mathbf{F}^*\mathbf{A}'\mathbf{F} &= \mathbf{F}^*\mathbf{A}\mathbf{F} - \mathbf{F}^*\mathbf{b}\mathbf{c}^T\mathbf{F} \\ &= \mathbf{D} - \mathbf{F}^*(\mathbf{b}\mathbf{c}^T\mathbf{F})\end{aligned}$$

where $\mathbf{D}$ is a diagonal matrix, and the term within parentheses is an $N \times N$ matrix having nonnull entries only in position (1, 1) and (2, 1). Moreover, these two entries have opposite sign. It turns out that $\mathbf{F}^*\mathbf{b}\mathbf{c}^T\mathbf{F}$ has nonnull elements only on the first column under the diagonal. This means that the matrix $\mathbf{A}'$ can be triangularized by means of the DFT matrix, and its eigenvalues (found on the diagonal) are the same as those of $\mathbf{A}$. This argument works for any number of oppositely signed couples of distinct values arbitrarily distributed in the vector $\mathbf{b}$. The same argument can be followed for proving the claim relative to the columns. In this case, we would start by forming the product $\mathbf{F}^*\mathbf{b}\mathbf{c}^T$. □

Note that the zeroth eigenvalue is no longer a "dc" eigenvalue. The corresponding eigenvector must be found in $\ker(\mathbf{A} - \mathbf{b}\mathbf{c}^T - \lambda_0\mathbf{I})$, where $\ker()$ gives the nullspace of its argument. In the case of a real circulant matrix $\mathbf{A}$ with eigenvalues along the unit circle, we have that $\lambda_0 = 1$ (the sum of the elements of a row of $\mathbf{A}$ is 1).

With the above choice of $\mathbf{b}$ and $\mathbf{c}$ coefficients, we obtain a perfectly flat amplitude response for equal-length delay lines. However, this is degenerate since this is the condition for pole–zero cancellation. As we will show in Section V, when using slightly different delay lengths, a nearly flat response at low frequencies is obtained as a perturbation of the pole–zero cancellation configuration.

B. Computational Complexity

In an Nth-order FDN, the core computations consist of N updates of the delay lines and a matrix by vector

multiplication. The delay line operations can proceed in parallel. The matrix by vector multiplication requires, in general, $O(N^2)$ operations (multiplications and additions). If the matrices arise from the scattering coefficients of a waveguide junction, the computations reduce to $O(N)$. The same order of complexity is required by the normalized waveguide junction. For the special case of a junction of equal-impedance waveguides, the multiplications can be replaced by shifts when N is a power of 2 [9]. In all these efficient cases, the eigenvalues of the feedback matrix are constrained to be at $+1$ or -1. The circulant matrix offers a more general eigenvalue distribution. Moreover, the matrix by vector multiplication can be implemented very efficiently in hardware. This multiplication can be viewed as a circular convolution of the column vector with the first row of the matrix. Such a convolution can be performed, when N is a power of 2, using two FFT's (one of which can be precomputed), an elementwise product between two N vectors, and an inverse FFT. The complexity of this algorithm is $O[N \log(N)]$. It is easy to implement this matrix-vector product in VLSI by means of the butterfly or other hypercubic architectures [24]. These architectures allow computations of the FFT in $O[\log(N)]$ time steps, and the algorithm can be pipelined.

The parallel implementation of waveguide scattering matrices cannot be done in less than $O[\log(N)]$ time steps because of the scalar product that is involved in Householder reflections of any kind. Hence, in parallel implementations, we lose the advantage of waveguide scattering matrices over circulant matrices.

Moreover, we can use number-theoretic Fourier transforms in order to compute the circular convolution. Such transforms work over commutative rings and can be arranged in such a way that all multiplications are replaced by shifts. Since in the convolution we have both the direct and inverse transforms, the overall result remains correct.

The circulant structure of $\mathbf{A}$ is advantageous for purposes of real-time control as well. In the first place, the entire matrix is determined by one of its rows or columns, and no matter how a row or column is modified, as long as the rest of the matrix is modified accordingly to preserve the circulant structure, the matrix will have unit-modulus eigenvalues as needed for losslessness. Furthermore, the top row of a circulant matrix is obtained from its eigenvalues by means of an inverse DFT. Therefore, it is possible to efficiently generate a continuous family of circulant matrices by continuously varying the complex phases of the eigenvalues. Moreover, if the matrix-vector multiplication is implemented in the frequency domain, the inverse DFT is not needed. Thus, we may move the N eigenvalues to arbitrary points on the unit circle and generate a wide family of efficiently computed lossless feedback matrices.

V. Applications

We have been using circulant networks for various purposes in sound synthesis and processing. Artificial reverberation is probably the most significant application, but other significant areas of interest can be found in sound synthesis and filtering.

A. Digital Reverberation

Two quantities have been proposed as criteria for measuring the "naturalness" of synthetic reverberation: the time density and the frequency density [7]. A good reverberator should provide high values of both densities, thus giving smooth, dense time and frequency responses.

The frequency density D_f is defined as the average number of resonances per Hertz. A general expression can be derived from the order of the system (5), assuming that all the poles are distinct, and no cancellation occurs:

$$D_f = \frac{1}{F_s} \sum_{i=1}^{N} m_i. \tag{32}$$

In real rooms, the frequency density increases at higher frequencies (as can be seen from (35) below).

In the prototype case, where the delay lines all have the same length m, we have

$$D_f = \frac{Nm}{F_s}. \tag{33}$$

The time density D_t is defined as the number of nonzero samples per second in the impulse response. In actual rooms, D_t is an increasing function of time. In order to obtain dense reverberation after the early reflections (e.g., after 80 ms), it helps to use different delay lengths.

The actual positions of frequency peaks depend on the feedback matrix and the delay lengths. If the delay lengths are fixed, we can vary some time-frequency properties of the structure simply by varying the distribution of eigenvalues of the feedback matrix. The total length of the delay lines should be chosen in such a way that the frequency density, as determined by (32), is high enough. Then, the matrix eigenvalues can be adjusted to avoid resonant peak clustering or other undesirable mode distributions.

It is interesting to discuss the effect of eigenvalues in the prototype case of equal delays. A uniform distribution of eigenvalues along the unit circle is optimum for the frequency response in the sense that it minimizes the maximum distance between peaks. However, it produces a highly repetitive time response. Conversely, clustering the eigenvalues around a point on the unit circle can be good for maximizing the length of time patterns, but the clustering of frequency peaks produces a poor reverberator amplitude response versus frequency. We see from these considerations that there is a time-frequency tradeoff. This tradeoff can be addressed using circulant matrices.

A couple of examples of different eigenvalue distributions are given in Fig. 3. The matrix $\mathbf{A}_2$ used in Fig. 3(b) is simply obtained by a right circular shift of the rows of the matrix $\mathbf{A}_1$, which is given by the junction of equal-impedance waveguides and, as already stated, has eigenvalues only at 1 and -1. We can express $\mathbf{A}_2$ as the product $\mathbf{A}_1\mathbf{\Pi}$, where $\mathbf{\Pi}$ is the right-shift matrix

$$\mathbf{\Pi} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}. \tag{34}$$

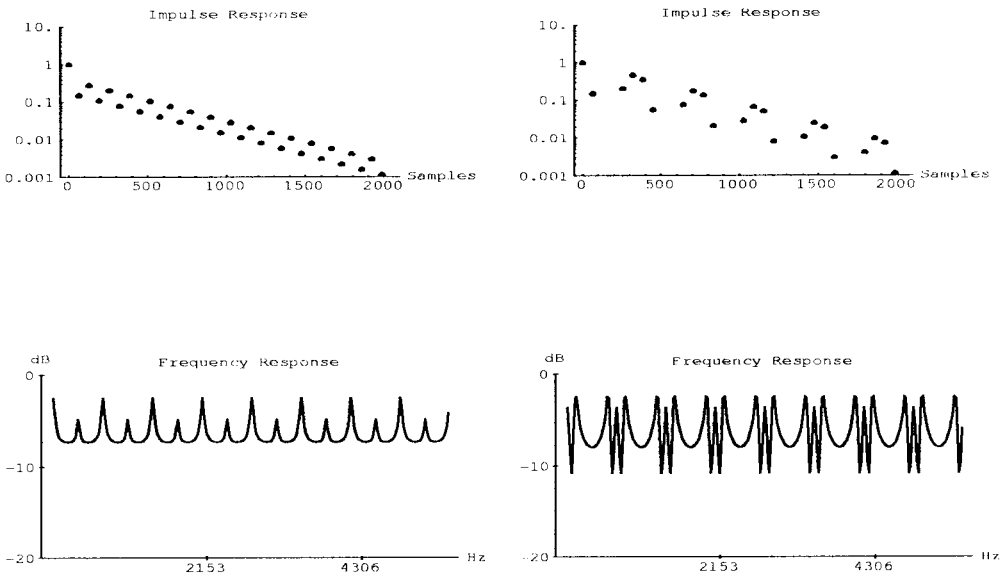

$$\mathbf{A}_1 = \begin{bmatrix} -1/3 & 2/3 & 2/3 \\ 2/3 & -1/3 & 2/3 \\ 2/3 & 2/3 & -1/3 \end{bmatrix} \qquad \mathbf{A}_2 = \begin{bmatrix} 2/3 & -1/3 & 2/3 \\ 2/3 & 2/3 & -1/3 \\ -1/3 & 2/3 & 2/3 \end{bmatrix}$$

$$\lambda_1 = \begin{bmatrix} -1 & 1 & -1 \end{bmatrix} \qquad \lambda_2 = \begin{bmatrix} e^{-j\pi/3} & 1 & e^{j\pi/3} \end{bmatrix}$$

(a) (b)

Fig. 3. Time and frequency behaviors for two circulant feedback delay networks that differ only by a shift on the rows of the feedback matrix.

Both $\mathbf{A}_1$ and $\mathbf{\Pi}$ are circulant, and therefore, the eigenvalues of $\mathbf{A}_2$ are given by the collection of the element-wise products of the eigenvalues of $\mathbf{A}_1$ and the eigenvalues of $\mathbf{\Pi}$, which are the Nth complex roots of 1 [23]. For clarity, we set all the delay lengths equal in the examples.

As a side comment, we notice that $\mathbf{\Pi}$ is the scattering matrix of the *circulator*, which is a circuit device that can be used to obtain the multiplication of one-port scattering parameters [25].

The shape of the frequency response depends also on the zeros, which were discussed in Section IV. In particular, Theorem 2 provides a way of setting the zeros exactly over the poles in the prototype equal-delay case. We anticipated in Section IV that the way to choose the vectors $\mathbf{b}$ and $\mathbf{c}$ indicated in Theorem 2 can be useful for getting a flat amplitude response at low frequencies when the delay lengths are slightly varied from the prototype case. Fig. 4 depicts the time and frequency responses for the CFDN using the same feedback matrix as in Fig. 3(b), having $\mathbf{b}^T = [1, 1, 1]$, $\mathbf{c}^T = [0, -1, 1]$, and delay lengths $m = [16, 17, 15]$. As we can see from Fig. 4, we are able to get a nearly flat amplitude response at low frequencies without losing the reverberating character of the time response. We believe that this is a good alternative to allpass filters, which tend to have degenerate impulse responses when the poles approach the unit circle.

B. Physical Room Modeling with FDN's

A flat amplitude response at low frequencies, while desirable in several practical situations, is not found in actual rooms. Therefore, if the goal is to model the reverberation of a physical room, the way indicated by Theorem 2 is not appropriate. Somewhat happily, the FDN can be the kernel of a model of rectangular room, and its parameters can be interpreted in a physical and geometrical framework. In this section, we give only a sketch of this framework since the details of the underlying metaphor are beyond the scope of this paper and can be found in [26].

Consider a lossless shoe-box shaped room having length l_x, depth l_y, and height l_z. For such a room, it is possible to compute analytically the frequencies of the normal modes [15] as

$$f_{n_x,n_y,n_z} = \frac{c}{2}\left[\left(\frac{n_x}{l_x}\right)^2 + \left(\frac{n_y}{l_y}\right)^2 + \left(\frac{n_z}{l_z}\right)^2\right]^{1/2} \tag{35}$$

where $n_x, n_y, n_z = 0, 1, 2, \cdots$, and c is the speed of sound in air. Each normal mode is associated with a direction in space, whose cosines, which are made by the wave propagation with

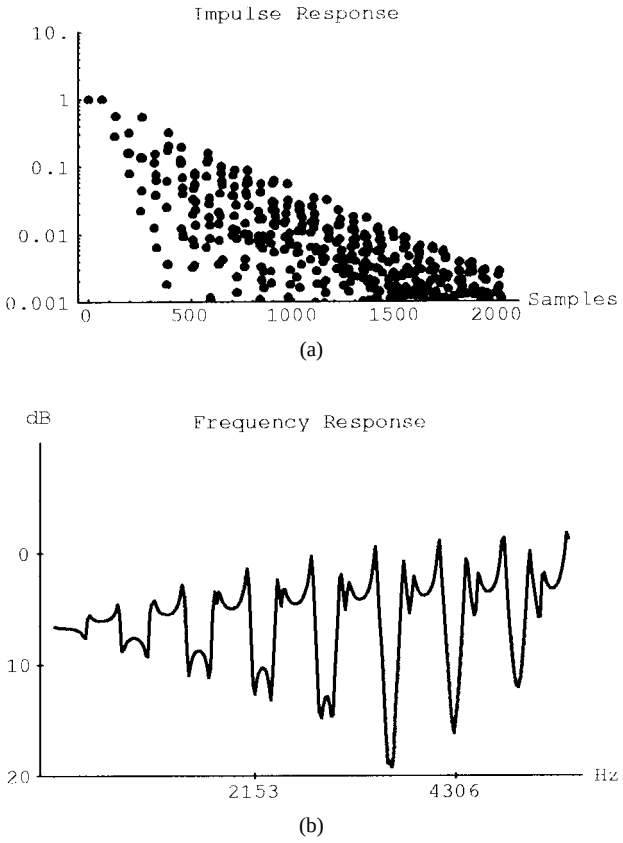

Fig. 4. Zero positioning that gives a nearly flat low-frequency response for the CFDN of Fig. 3(b), with $\mathbf{b} = [1 \quad 1 \quad 1]$, $\mathbf{c} = [0 \quad -1 \quad 1]$, and $Delays = [16 \quad 17 \quad 15]$.

respect to the x, y, and z axes, are

$$\begin{aligned} \frac{v_x}{v} &= \frac{n_x c}{2fl_x} \\ \frac{v_y}{v} &= \frac{n_y c}{2fl_y} \\ \frac{v_z}{v} &= \frac{n_z c}{2fl_z} \end{aligned} \tag{36}$$

where v is the magnitude of the (vector) spatial frequency, and the subscripts in f_{n_x, n_y, n_z} have been dropped for conciseness.

The triple $\mathbf{n} = (n_x, n_y, n_z)$ completely characterizes a normal mode. All the triples that are multiples are associated with a harmonic series of frequencies and with the same direction in space. This suggests that any harmonic series of normal mode frequencies can be obtained by means of a linear resonator (in other words, a comb filter) whose length in seconds is set to $d_0 = 1/f_0$, where f_0 is the fundamental frequency of the harmonic series. Therefore, we can decompose the modal distribution of the response of an actual room into harmonic subsets (a harmonic of f_{n_x, n_y, n_z} is obtained by multiplying n_x, n_y, and n_z by the same integer). Sorting these harmonic subsets according to their fundamental frequencies and taking the reciprocals of the N lowest fundamental frequencies yields a parallel comb filter representation of the room (i.e., an FDN with diagonal feedback matrix) so that the FDN reproduces the lowest eigenfrequencies exactly. This procedure was already outlined in [7] as a mean of identifying the parallel comb-filter parameters from a measured impulse response.

We can elaborate the representation further by interpreting the quantity d_0 as the time taken by a plane wavefront to travel a certain distance along the direction (36) in space. In fact, a normal mode and all its harmonically related multiples can be thought of as a planewave bouncing back and forth in the closed environment [15]. For a finite medium, in order to support such an infinite plane wave, the planar fronts have to be bent at the walls such that they form a constant-area closed surface. It can be verified that the time d_0 is the time interval between two successive collisions of two plane wavefronts.

Once it is established that in an idealized rectangular room each harmonic subset of normal modes can be represented by a linear resonator oriented along a given direction in space, we can introduce other "second-order" effects into the basic model.

Let us consider an octant in space. Taking the first N fundamental frequencies in the harmonic-subset decomposition of the normal modes corresponds to sampling in space along N directions. An object in any point of the space will provide scattering among the N directions. The walls themselves, when they are not ideally smooth, scatter the waves in different directions. We can think of lumping all these diffusion effects and representing them in the nondiagonal elements of the scattering matrix of a FDN. With some approximation, it was also shown in [26] that an isotropic object in a nondiffusive rectangular room can be represented by a circulant matrix, provided that the spatial sampling is almost uniform, and the proper ordering of directions is chosen.

The geometric interpretation allows one to properly excite the modes according to the position of the sound source by simply replacing each coefficient in the vector $\mathbf{b}$ with a suitable cascade of FIR comb filters [26]. The position at which we listen to the sound is related to the $\mathbf{c}$ coefficients in a similar way. It is also quite easy to take into account the radiation pattern of the source and the directivity of the pick up. Perhaps more importantly, the absorption coefficients of the walls can be made to be direction dependent, as they are found in reality, because they affect the different "linear resonators" differently.

In the model at hand, the matrix element $a_{i,j}$ scales the signal transmission from mode j to mode i. The diagonal of the feedback matrix determines the strength of the "standing waves" set up along each pattern. Equivalently, we can think of a DWN modeling the parallel junction of N acoustic tubes, where each tube gives rise to a harmonic subset of normal modes.

The physical modeling viewpoint is limited by the fact that only N "standing-wave paths" in the room are being simulated, and all nonspecular reflections are being forced to enter some subset of the supported ray paths.

In the model, the diffusivity of the whole reverberation is lumped in the properties of the scattering matrix. This is a dramatic simplification, but it allows better control of diffusivity in isolation from other room parameters.

The geometrical interpretation is useful for computing the lengths of the delay lines according to the dimensions of a particular room since each wavefront path corresponds to a

normal mode. In previous work on artificial reverberation [1] and [2], the choice of the delay-line lengths in the allpass and combfilter sections is a primary issue. Typically, the choice is guided by heuristic rules or number-theoretic criteria, and a lot of trial and error is often necessary to obtain good values.

C. Physical Room Modeling with DWN's

Recent developments in physical modeling using digital waveguides have included the use of a *waveguide mesh* to model 2-D membranes and 3-D rooms [27]–[29]. In the membrane, for example, a rectilinear mesh of digital waveguides can be interconnected via four-port scattering junctions to provide lossless prototypes for "plate reverberators" and the like. A single dispersive waveguide (made dispersive using embedded allpass filters) can be used to model "spring reverberators." Savioja *et al.* [29] have found that the rectilinear 3-D waveguide mesh has good room simulation properties at low frequencies.

Since reverberation quality generally increases with the number dimensions (from spring to plate to acoustic space), it is plausible to expect that higher dimensional waveguide meshes will provide better reverberation than we have ever known. Generalizing (35) to higher dimensions, one can see that the higher the dimensionality, the more rapidly the mode density increases with frequency. However, there exists a "Schroeder limit" at which the average spacing between eigenfrequencies becomes substantially smaller than the bandwidth of one mode [30]. Above this limit, there is no reason to increase the frequency density since it will bring no audible improvement [7]. Nevertheless, it would be interesting to pursue the study and implementation of high dimensional waveguide meshes.

The waveguide mesh is structurally lossless so that there is no attenuation error in the sampled wave propagation. However, the grid quantization gives rise to *dispersion* error: The speed of sound effectively varies somewhat as a function of frequency and propagation direction on the mesh. Generally, results are very accurate at low frequencies, but sound speed decreases gradually as frequency increases in all but certain directions that tend to be diagonals along the mesh [27]. The choice of mesh geometry has a strong effect on the dispersion behavior [28]. It also strongly affects computational complexity. As an example, whenever an isotropic mesh utilizes N-port scattering junctions in which N is a power of 2, the scattering matrices require no multiplies [9]. For rectilinear meshes, membranes are multiply free, as are solids in 4-D (since the number of ports is $2N$, in N-dimensional space). The tetrahedral mesh, which is analogous to the diamond crystal, requires no multiplies to fill 3-D space. Multiply-free waveguide meshes can be integrated very densely in VLSI.

A final word about waveguide meshes is that they, like any other LTI systems, can be expressed in a sparse state-space form that yields an FDN that can be interpreted as a physical model.

D. Practical FDN Design

In our experience, given an FDN reverberator structure, setting the delay lengths can be a rather tedious job. The vast majority of possible delays provide poor results in the sense that the time response is too "rough" or the frequency response is too "colored." An interesting approach to this problem might be to use nonlinear optimization techniques such as "simulated annealing" or "genetic algorithms" to optimize the delay lengths such that "perceptual uniformity" of the response is maximized in the time and frequency domains jointly.

Designing the delay lengths from room geometry has the property of giving a reverberator that is always consistent with a desired room in that the low-frequency modes are matched. However, there does not seem to be any compelling reason to match specific low-frequency mode tunings. Noticeable room resonances are normally perceived as defects in a listening space. Early reflections, on the other hand, contribute strongly to the perceived "spatial impression" [31]. In other applications, however, such as modeling the soundboard of a piano as a reverberator, the specific coloring or "equalization" provided by the reverberator is important and must be preserved. In such applications, it is normally necessary to match low-frequency resonances accurately and high-frequency resonances only statistically.

When the FDN order is large (larger than eight for satisfactory results), poor results can still be obtained when modeling desired room dimensions that are not favorable. In fact, even for the shoe-box room shape, the relative dimensions play a very important role in determining the smoothness of the reverb [32]. Of course, *diffusion* contributes significant smoothing to the response; therefore, full feedback matrices (as opposed to diagonal feedback matrices) are especially needed to achieve good reverberators using low-order FDN's.

On balance, it seems that what is needed for good reverberator design, in general, is the following:

1) precise matching of early reflections,
2) minimal coloration due to uneven mode distributions in the frequency domain,
3) an appropriate smoothly declining decay-time versus frequency,
4) smooth, rich echo density late in the impulse response having no noticeable patterns.

These desiderata indicate that rather than attempting to model real rooms, lossless prototype FDN's optimizing criteria (2) and (4) should be found, for a given order, which have at least one delay line long enough to support injection of specific early reflections to satisfy (1), and then, lowpass filters as in (21) should be added to satisfy criterion (3). The main open issue is how the optimization of (2) and (4) should best be carried out for specific classes of structurally lossless feedback matrices.

E. Resonators

FDN's with short delay lines may be used to produce resonances irregularly spread over frequency. A possible application could be the simulation of resonances in the body or soundboard of a string instrument.

Mathews and Kohut [33] showed that in this kind of simulation of the violin body, the exact position and height of resonances is not usually important; on the contrary, they

stated that the Q's of the resonances must be sufficiently large and the peaks must be sufficiently close together. Thus, even in rather small physical resonators, a statistical matching may be as effective as a more precise, mode-for-mode matching.

With CFDN's, we can easily achieve these goals, and we can vary the distribution of peaks by acting on the delay lengths and/or the feedback matrix. In this context, the main advantage of using CFDN's over general FDN's is that the feedback matrix has N parameters that are related to the eigenvalues by means of a DFT. This means that we have the possibility of controlling a large number of resonances using a number of parameters that are linear in the order of the structure, and the parameter control can be performed very efficiently and safely, i.e., without running into instabilities. The control over matrix eigenvalues is complementary with respect to the control of delay lengths: While changing a delay has a stretching or squeezing effect on resonance positions all along the frequency axis, changing the eigenvalues produces alternative changes in the distribution of resonances, such as clustering the peaks, as is illustrated in Fig. 3.

Another interesting application of CFDN's is as resonators in a feedback loop for pseudo-physical sound-synthesis techniques. By exciting these structures with bursts of white noise, we obtain a multivariable extension of the Karplus–Strong algorithm [34] that is very effective for simulating membranes and bars. Alternatively, we can couple these resonators with nonlinear exciters and explore new families of sustained sounds, as in waveguide synthesis [14].

We have been using effectively CFDN's in live electronic performances, where the exciting signal is coming from a traditional instrument, and the CFDN provides a complicated filtering pattern whose frequency shape can be controlled in real time by its parameters (eigenvalues or row elements).

VI. Conclusion

This paper presented generalizations and new special cases for the matrix used in feedback delay networks. In particular, necessary and sufficient conditions were derived for losslessness of such a matrix. The correspondence between FDN's and digital waveguide networks can be used to obtain FDN parameters based on the physics and geometry of a real acoustic space, rather than by rules of thumb or number-theoretic rules.

In proposing the CFDN structure, we have tried to achieve two goals: efficiency and versatility with respect to the time-frequency behavior. Efficiency is achieved by taking advantage of the circulant structure of the feedback matrix, and it increases with the size of the matrix. Versatility is achieved by introducing the matrix eigenvalues into the design process for artificial reverberators. Passing from the eigenvalues to the matrix coefficients requires only a single inverse DFT or FFT. Eigenvalues act on the distribution of frequency peaks, thus giving controls pertaining to the color and smoothness of the reverberation.

In addition to application of CFDN's in artificial reverberation, we have outlined some other uses as resonators in sound synthesis and processing.

Acknowledgment

The authors wish to thank the anonymous reviewers for their excellent comments and suggestions.

References

[1] M. R. Schroeder, "Natural-sounding artificial reverberation," *J. Audio Eng. Soc.*, vol. 10, no. 3, pp. 219–233, 1962.
[2] J. A. Moorer, "About this reverberation business," *Comput. Music J.*, vol. 3, no. 2, pp. 13–18, 1979.
[3] M. A. Gerzon, "Unitary (energy preserving) multichannel networks with feedbacks," *Electron. Lett.*, vol. 12, no. 11, pp. 278–279, 1976.
[4] P. P. Vaidyanathan, *Multirate Systems and Filter Banks.* Englewood Cliffs, NJ: Prentice-Hall, 1993.
[5] J. Stautner and M. Puckette, "Designing multichannel reverberators," *Comput. Music J.*, vol. 6, no. 1, pp. 52–65, Spring 1982.
[6] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing.* Englewood Cliffs, NJ: Prentice-Hall, 1975.
[7] J.-M. Jot, "Etude et reálization d'un spatialisateur de sons par modèles physiques et perceptifs," Ph.D. thesis, Telecom, Paris, 92 E 019, 1992.
[8] J.-M Jot and A. Chaigne, "Digital delay networks for designing artificial reverberators," in *Audio Eng. Soc. Conv.*, Feb. 1991.
[9] J. O. Smith, "A new approach to digital reverberation using closed waveguide networks," in *Proc. 1985 Int. Comput. Music Conf.*, Vancouver, Computer Music Assoc., 1985, pp. 47–53. Also available in [16].
[10] ———, "Elimination of limit cycles and overflow oscillations in time-varying lattice and ladder digital filters," CCRMA, Music Dept., Stanford Univ., Tech. Rep. STAN-M-35, May 1986, Short version published in *Proc. IEEE Conf. Circuits Syst.*, San Jose, May 1986. Full version also available in [16].
[11] J. M. Ortega, *Numerical Analysis.* New York: Academic, 1972.
[12] T. Kailath, *Linear Systems.* Englewood Cliffs, NJ: Prentice-Hall, 1980.
[13] D. A. Jaffe and J. O. Smith, "Extensions of the Karplus–Strong plucked string algorithm," *Comp. Music J.*, vol. 7, no. 2, pp. 56–69, 1983.
[14] J. O. Smith, "Physical modeling using digital waveguides," *Comput. Music J.*, vol. 16, no. 4, pp. 74–91, Winter 1992, Special Issue: *Physical Modeling of Musical Instruments, Part I.*
[15] P. M. Morse, *Vibration and Sound*, Amer. Inst. of Phys. for the Acoust. Soc. Amer., 1976, first ed. 1936, second ed. 1948).
[16] J. O. Smith, "Music applications of digital waveguides," CCRMA, Music Dept., Stanford Univ., Tech. Rep. STAN-M-39, 1987, A compendium containing four related papers and presentation overheads on digital waveguide reverberation, synthesis, and filtering. CCRMA Tech. Reps. can be ordered by calling (415) 723-4971 or by sending an e-mail request to hmk@ccrma.stanford.edu.
[17] J. D. Markel and A. H. Gray, *Linear Prediction of Speech.* New York: Springer, 1976.
[18] G. H. Golub and C. F. Van Loan, *Matrix Computations.* Baltimore: The Johns Hopkins University Press, 1989.
[19] V. Belevitch, *Classical Network Theory.* San Francisco: Holden Day, 1968.
[20] S. Perlis, *Theory of Matrices.* New York: Dover, 1991, first ed., 1952.
[21] J. O. Smith and D. Rocchesso, "Connections between feedback delay networks and waveguide networks for digital reverberation," in *Proc. 1994 Int. Comp. Music Conf.*, Århus, Comput. Music Assoc., 1995, pp. 376–377.
[22] A. H. Gray, "Passive cascaded lattice digital filters," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-27, no. 5, pp. 337–344, May 1980.
[23] P. J. Davis, *Circulant Matrices.* New York: Wiley, 1979.
[24] F. T. Leighton, *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes.* Los Altos, CA: William Kaufmann, 1992.
[25] R. W. Newcomb, *Linear Multiport Synthesis.* New York: McGraw-Hill, 1966.
[26] D. Rocchesso, "The ball within the box: A sound-processing metaphor," *Comput. Music J.*, vol. 19, no. 4, pp. 47–57, Winter 1995.
[27] S. A. Van Duyne and J. O. Smith, "Physical modeling with the 2-D digital waveguide mesh," in *Proc. 1993 Int. Comput. Music Conf.*, Tokyo, Comput. Music Assoc., 1993, pp. 40–47.
[28] ———, "The tetrahedral waveguide mesh: Multiply-free computation of wave propagation in free space," in *Proc. IEEE Workshop Applications Signal Processing Audio Acoust.*, New Paltz, NY, Oct. 1995.
[29] L. Savioja, J. Backman, A. Järvinen, and T. Takala, "Waveguide mesh method for low-frequency simulation of room acoustics," in *Proc. 15th Int. Conf. Acoust. (ICA-95)*, Trondheim, Norway, June 1995, pp. 637–640.

[30] M. R. Schroeder, "Statistical parameters of the frequency response curves of large rooms," *J. Audio Eng. Soc.*, vol. 35, no. 5, pp. 299–306, May 1987.
[31] J. Borish, "Electronic simulation of auditorium acoustics," Ph.D. thesis, Elec. Eng. Dept., Stanford Univ., 1984; available as CCRMA Tech. Rep. STAN-M-18.
[32] J. Milner and R. Bernhard, "An investigation of the modal characteristics of nonrectangular reverberation rooms," *J. Acoust. Soc. Amer.*, vol. 85, no. 2, pp. 772–779, 1989.
[33] M. Mathrews and J. Kohut, "Electronic simulation of violin resonances," *J. Acoust. Soc. Amer.*, vol. 53, no. 6, pp. 1620–1626, 1973.
[34] K. Karplus and A. Strong, "Digital synthesis of plucked string and drum timbres," *Comput. Music J.*, vol. 7, no. 2, pp. 43–55, 1983.

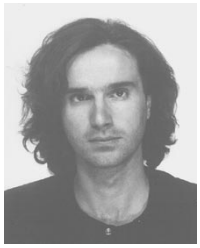

Davide Rocchesso received the *Laurea in Ingegneria Elettronica* degree from the Università di Padova in 1992 with a thesis on real-time physical modeling of musical instruments. He is currently a Ph.D. candidate at the Dipartimento di Elettronica e Informatica, Università di Padova, Italy.

In 1994 and 1995, he was a visiting scholar at the Center for Computer Research in Music and Acoustics (CCRMA), Stanford University, Stanford, CA. Since 1991, he has been collaborating with the Centro di Sonologia Computazionale (CSC) dell'Università di Padova as a researcher and a live-electronic designer/performer. His main interests are in audio signal processing, physical modeling, sound reverberation and spatialization, parallel algorithms. Since 1995, he has been a member of the Board of Directors of the Associazione di Informatica Musicale Italiana (AIMI). His home page on the Web is http://www.dei.unipd.it/english/csc/roc/intro.html.

Julius O. Smith received the B.S.E.E. degree from Rice University, Houston, TX, in 1975. He received the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1978 and 1983, respectively. His Ph.D. research involved the application of digital signal processing and system identification techniques to the modeling and synthesis of the violin, clarinet, reverberant spaces, and other musical systems.

From 1975 to 1977, he worked in the Signal Processing Department at ESL, Sunnyvale, CA, on systems for digital communications. From 1982 to 1986, he was with the Adaptive Systems Department at Systems Control Technology, Palo Alto, CA, where he worked in the areas of adaptive filtering and spectral estimation. From 1986 to 1991, he was employed at NeXT Computer, Inc., where he was responsible for sound, music, and signal processing software for the NeXT computer workstation. Since then, he has been an Associate Professor at the Center for Computer Research in Music and Acoustics (CCRMA) at Stanford teaching courses in signal processing and music technology and pursuing research in signal processing techniques applied to music and audio. His home page on the Web is http://www-ccrma.stanford.edu/jos/.