Exploration of Quantum Feedback Delay Networks

Davide Rocchesso[1]

[1]Department of Mathematics and Computer Science – University of Palermo

davide.rocchesso@unipa.it

Abstract. *Quantum Feedback Delay Networks are audio-processing structures based on delay lines and scattering matrices in a feedback loop, where audio is encoded in qubits, and qubits or their representations evolve through the network. In this exploration, different realizations with different degrees of physical realizability are tried out, starting from the quantum version of a recursive comb filter, up to higher order structures with several delay lines and qubits.*

1 Introduction

Feedback Delay Networks (FDN) are structures that found extensive use in artificial reverberation and digital audio effects [1, 2]. They can be seen as an extension of the recursive comb filter, where the single delay line is replaced by a battery of delay lines, and the feedback coefficient is replaced by a feedback matrix, possibly accompanied by per-line coefficients and filters. Audio samples are injected into the FDN delay lines as input, and a weighted sum of the outputs of delay lines is taken as an instantaneous output audio sample. A geometric interpretation was proposed of the FDN as a scattering object (the matrix) within a box with reflecting walls [3]. FDNs are designed starting from lossless prototypes, and then shaping their time-frequency characteristics based on some desired properties, such as densities of echoes and resonances, through tuning of delay lines and insertion of coefficients and filters. The feedback matrices that can be used for lossless prototypes have been thoroughly studied [4, 5] and, among these, unitary matrices have found large use. This simple fact makes the FDN structure an attractive playground to experiment with quantum computing tools, at the simulation level as well as with physical realizations, whenever these are possible.

Among the possible ways to encode audio as quantum bits, Quantum Pulse Audio Modulation (QPAM) [6] seems to be suitable for sample-by-sample audio processing, as it encodes a discrete audio sample as a probability amplitude of a quantum state. Although in QPAM the set of 2^n probability amplitudes of a n-qubit quantum state refers to an amplitude-normalized segment of 2^n audio samples, for sample-by-sample audio processing we may consider all of these numbers to refer to the same time instant, as if they were samples of different audio channels.

This paper explores[1] the idea of applying the FDN structure to audio-encoding qubits instead of discrete-time audio signals. If unitary matrices are used, the scattering operation can be read as a unitary evolution of the quantum state found at the output of delay lines.

Our exploration of QFDNs starts from the simplest structure, involving one-qubit feedback, that is the quantum analog of the recursive comb filter. Then the exploration addresses multi-qubit states. Several implementation choices can be made, that make the QFDN realization more or less amenable for direct implementation on a quantum computer. Since measurement and entanglement can be introduced in the structure, the behavior of the QFDN turns out to be very different from the linear response of the classical FDN.

[1]The reported examples are available, as a jupyter notebook containing the full code, on https://github.com/d-rocchesso/QFDN

2 The recursive Quantum Comb Filter

The recursive comb filter is a basic ingredient for fundamental sound synthesis (as in the Karplus-Strong algorithm [7]) and artificial reverberation [8, 9]. Essentially, the comb filter is a feedback loop containing a delay line, a multiply/filter for in-loop attenuation, and an adder for input signal. The name comb comes from the shape of its magnitude frequency response, that emphasizes those frequencies that are multiples of a fundamental, like the teeth of a comb. With exploratory attitude, we try to understand what happens if we extend a similar structure to digital audio encoded as qubits, being aware that the behavior may be highly different, and mostly nonlinear, especially when we deviate from physical realizability or we introduce measurement in the loop.

2.1 One delay

Figure 1 shows the structure of the recursive Quantum Comb Filter (QCF), relying on a single delay line and on a feedback matrix, that we call the Combgate, as it can be realized as a single-qubit gate.

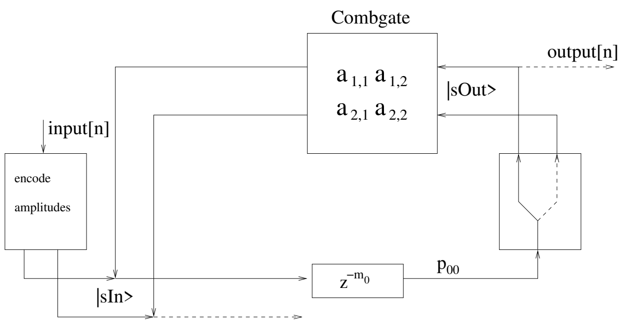

Figure 1: The recursive Quantum Comb Filter

2.1.1 Amplitude encoding

As in QPAM, instantaneous audio input may be scaled and shifted to stay between 0 and 1, and set to represent the first probability amplitude of a qubit. The zero signal value is set at probability-amplitude value $1/\sqrt{2}$, and the minus-one signal is set to 0. If we consider a state $[1,0]'$, corresponding to signal value one, and evolve it through a Hadamard gate, we get the state $\left[1/\sqrt{2}, 1/\sqrt{2}\right]'$, that would correspond to a zero signal value. Similarly, the state $[0,1]'$ would be Hadamard-evolved to $\left[1/\sqrt{2}, -1/\sqrt{2}\right]'$, that would be another form of zero signal.

2.1.2 Delay lines

In classical digital audio, delay lines are simply implemented as circular buffers, and fractional time delays can be achieved through some form of interpolation [10, 11]. In quantum digital audio, qubits are not easily delayed nor managed in circular buffers. However, in a simulation environment, we can think of knowing and delaying one of the two probability amplitudes defining the state of a qubit, knowing that the other must be power-complementary. Actually, there are infinitely many possibilities for the power-complementary probability amplitude, but we may choose the one that is real and positive.

If delays carry probability amplitudes, we initialize them in such a way that an identity feedback matrix would produce a constant output identical to the zero input. This requires initializing the delays with value $1/\sqrt{2}$.

2.1.3 Summation nodes

The converging arrows of figure 1 represent summation nodes, or points where two valid quantum states produce a valid quantum state by some kind of summation.

Having a probability-encoding of the input that produces a legitimate state vector

$$|\psi\rangle = \left[\sqrt{input}, \sqrt{1 - input}\right]' = a_0\,|0\rangle + b_0\,|1\rangle\,,$$

we can sum it to the state that results for matrix-evolution of $|sOut\rangle$, that is

$$|\phi\rangle = a_1\,|0\rangle + b_1\,|1\rangle\,,$$

as

$$|\psi\rangle + |\phi\rangle = \sqrt{\frac{|a_0|^2 + |a_1|^2}{2}}\,|0\rangle + \sqrt{\frac{|b_0|^2 + |b_1|^2}{2}}\,|1\rangle\,.$$

This operation forces the probability amplitudes at the entrance of delay lines to be real valued and positive, and to form a valid quantum state.

2.1.4 Identity feedback matrix

If an identity matrix is used for matrix feedback, a constant zero audio input produces a constant zero audio output. On the other hand, a unit impulse in the input produces a decaying impulse train, with amplitudes $0.75, 0.625, 0.5625, 0.53125, \ldots$. Subtracting the 0.5 offset we get a perfect division by two at every cycle, as in a classical recursive comb filter with loop coefficient set at 0.5. The response to a damped sinusoid is a distorted and delay-modulated damped sinusoid, as depicted in figure 2.

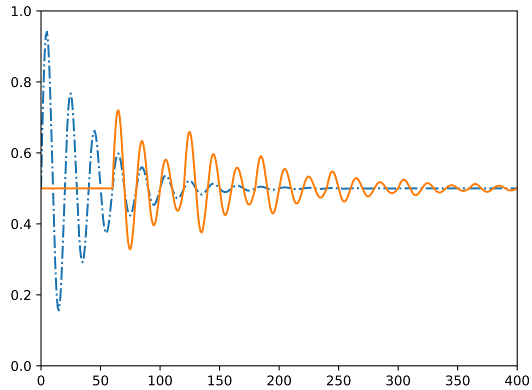

Figure 2: Response (continuous line) of the 1-qubit QCF with identity feedback matrix to a damped sinusoid (dash-dotted line). Time is measured in samples and delay is set equal to 59 samples.

2.1.5 Hadamard feedback matrix

With a Hadamard feedback matrix the behavior is quite far from that of classical recursive comb filters. A constant zero input, with delay initialized at $1/\sqrt{2}$, generates a staircase oscillation with peak at 0.75, and converging to 0.7236. A unit impulse in the input produces decaying impulses alternating their signs around the staircase oscillation. A damped sinusoid gets heavily distorted as an effect of repeated circulation in the loop (see figure 3).

In the Bloch sphere, the Hadamard operator is equivalent to a rotation by $-\pi/2$ around the y axis, followed by a phase flip (Z gate). With rotations between 0 and $-\pi/2$ about the y axis, we can get behaviors that are intermediate between identity and Hadamard, in terms of decay time and asymptotic value. For example, for a rotation of $\pi/16$, the response to a damped sinusoid is depicted in figure 4, and can be compared to figures 2 and 3.

2.2 Two delays

In the realization of figure 1, the only part that is amenable to quantum computation is the Combgate, that is a single-qubit unitary operator. The states $|sIn\rangle$ and $|sOut\rangle$ must indeed be constructed out of classical operations. If a quantum delay line would be available, however, we may think of having the state $|sOut\rangle$ as a delayed version of $|sIn\rangle$. In simulation, a quantum delay line would be realized with two equal-length lines that delay the two probability amplitudes. The circuit is almost equivalent to that of the single delay, with the difference that the state phase (and sign) can be maintained through the loop. The encoded input gets mixed in-phase with the feedback qubit.

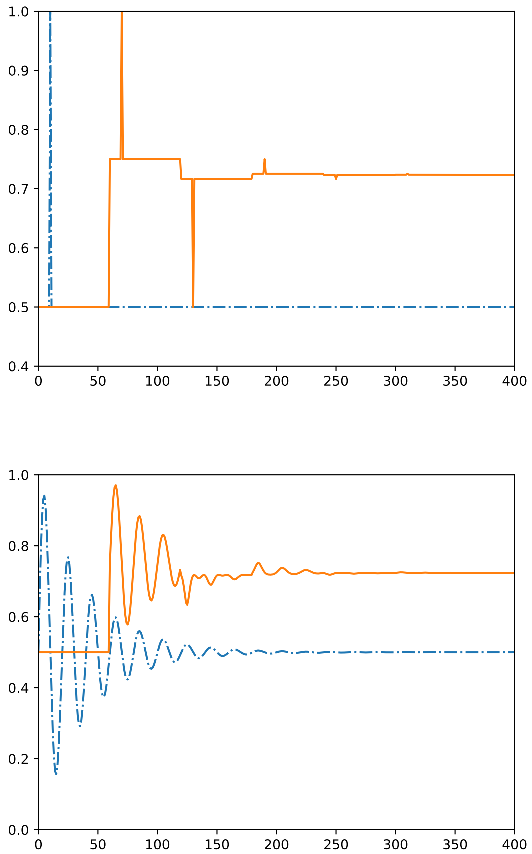

Figure 3: Response (continuous line) of the 1-qubit QCF with Hadamard feedback matrix to: (Top) a unit pulse and (Bottom) a damped sinusoid. Time is measured in samples and delay is set equal to 59 samples.

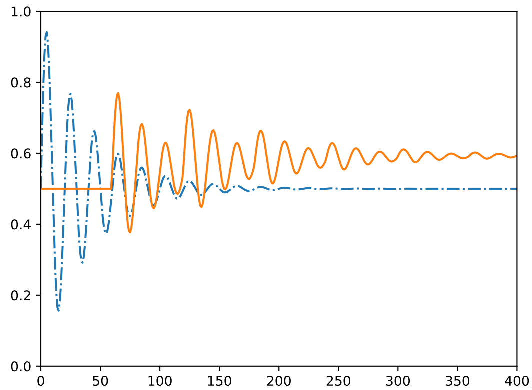

Figure 4: Response (continuous line) of the 1-qubit QCF with feedback matrix corresponding to a $-\pi/16$ rotation around the y axis, followed by a phase flip, to a damped sinusoid. Time is measured in samples and delay is set equal to 59 samples.

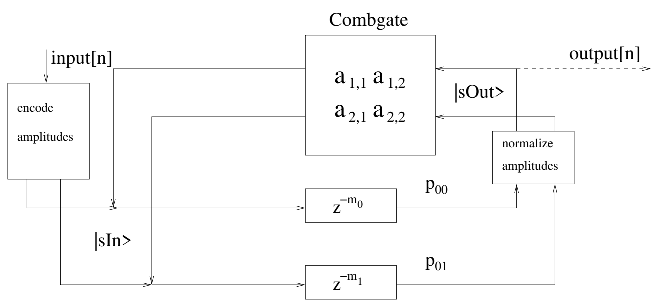

Figure 5: The recursive Quantum Comb Filter, with two delay lines, possibly of different length.

If we stick with simulation, the more general circuit of figure 5 is possible, where the two classical lines may take different values of delay. For such a general (non-physical) case, a state normalization stage has to be applied at the exit of the delay lines.

With identity or Hadamard combgate the behavior of the circuit with identical delay lengths is exactly as that of the circuit with a single delayed probability amplitude, described in sec. 2.1.

It is with rotations between 0 and $-\pi/2$ about the y axis, between identity and Hadamard matrices, that we get a peculiar oscillatory behavior, with a square wave actually modulating the response, as depicted in figure 6.

When excited by a vocal trill, the latter configuration gives a modified trill, with the square wave that tends to take over during silences, as depicted in figure 7.

2.3 Minimal QCF in state-space form

Consider the QCF with two equal-length delays, where the delay length is reduced to two samples. Excluding input and output, the feedback delay loop corresponds to repeated evolution of a two-qubit quantum state, which is specified by the four probability amplitudes stored in each of the delay units, represented by variables p_1, p_2, s_1, s_2 in figure 8.

The system state is

$$|w\rangle = [p_1, p_2, s_1, s_2]'$$

and the state-space-evolution matrix is

$$\mathbf{A}_{ss} = \begin{bmatrix} 0 & 0 & a_{1,1} & a_{1,2} \\ 0 & 0 & a_{2,1} & a_{2,2} \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \quad (1)$$

The two-qubit quantum state can be initialized to perfect superposition, with all probability amplitudes set to $1/\sqrt{4} = 1/2$, which is different from the previous case where the evolution is on a single qubit at a time.

To extract an output in QPAM according to the realization of figure 8, we extract the probability of the third element of vector $|w\rangle$, corresponding to number s_1, or probability of qubits being in state $[1, 0]'$.

To include an input, a tensor product may be formed between the state and the input, and the resulting

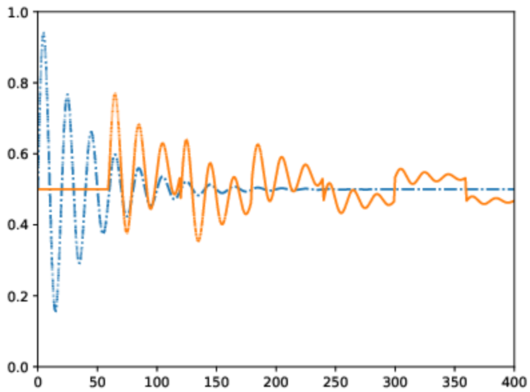

Figure 6: Response (continuous line) of the 2-delays 1-qubit QCF with feedback matrix corresponding to a $-\pi/16$ rotation around the y axis, followed by a phase flip, to a damped sinusoid. Time is measured in samples and delays are set equal to 59 samples.

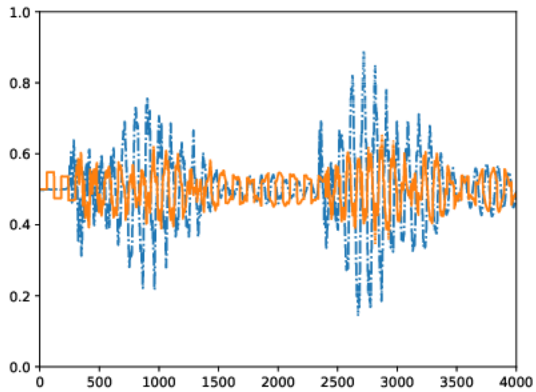

Figure 7: Response (continuous line) of the 2-delays 1-qubit QCF with feedback matrix corresponding to a $-\pi/16$ rotation around the y axis, followed by a phase flip, to a vocal trill. Time is measured in samples and delays are set equal to 59 samples.

state can be evolved by a 3-qubit circuit. If one of the qubits is traced-out, a new 2-qubit state is formed. In this way, either the matrix or the input turn out to be irrelevant. However, if we entangle (input) qubit 2 with qubit 0, as depicted in figure 9.top, and we measure qubit 2 before letting the two other qubits evolve, we get an interesting effect: with y-rotation and phase-flip we get a sort of micro-reverberation, with stochastic spikes having amplitude that can be controlled by the rotation phase. With the Combgate of figure 8 set to the identity matrix, a damped sinusoid gets transformed as in figure 9.bottom.

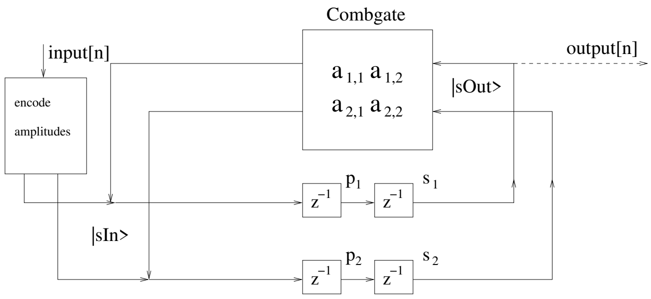

Figure 8: The recursive QCF with 2-sample delays

3 The Quantum Feedback Delay Network

We can extend the QCF to handle multiple qubits. A notable case is that of two qubits, with a 4×4 feedback matrix.

3.1 One delay per qubit

Figure 10 shows the structure of the QFDN, relying on a single delay line for each of the two qubits, and on a general feedback matrix, that we call the FDNgate, as it can be realized as a two-qubit gate. The delay lines carry probabilities, and blocks to convert amplitudes to and from probabilities are needed.

3.1.1 The feedback matrix

For the 4×4 size there exists the specific matrix (2) that is circulant, it is a Householder reflector, and has Hadamard property (i.e., it is made of ± 1 and has orthogonal rows). The matrix, interpreted as a scattering object, is maximally diffusive [12].

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

The matrix, as converted into a quantum circuit, evolves the state $|00\rangle$ into $-\frac{1}{2}|00\rangle + \frac{1}{2}|01\rangle + \frac{1}{2}|10\rangle + \frac{1}{2}|11\rangle$. In terms of rotations and c-nots, the matrix can be realized as in figure 11.

3.1.2 Amplitude encoding

As in QPAM, instantaneous audio input may be scaled and shifted to stay between 0 and 1, and set to represent the first probability amplitude of a qubit. The other probability amplitudes can be set equal and power complementary, i.e., the sum of the squares of all amplitude magnitudes gives one.

3.1.3 Summation nodes

The summation nodes, given

$$|\psi >= a_0|00> + b_0|01> + c_0|10> + d_0|11>$$

and

$$|\phi >= a_1|00> + b_1|01> + c_1|10> + d_1|11>,$$

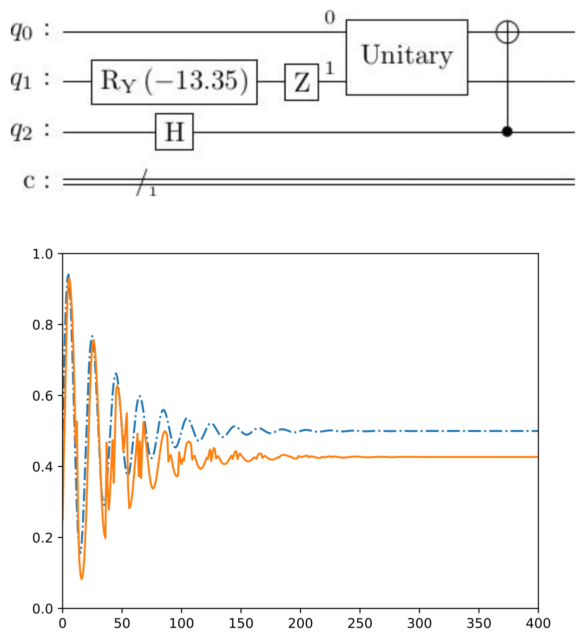

Figure 9: Top: Entangling the input (q_2) with the evolved state in state-space representation (the `Unitary` corresponds to A_{ss} of equation (1)); Bottom: Response (continuous line) to a damped sinusoid with identity feedback matrix, phase y-rotation and phase flip.

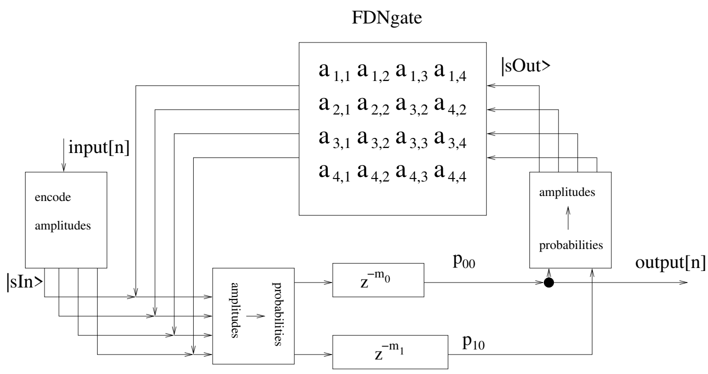

Figure 10: The QFDN for two qubits

compute

$$|\psi> + |\phi> = a_s|00> + b_s|01> + c_s|10> + d_s|11> = \sqrt{\frac{|a_0|^2 + |a_1|^2}{2}}|00> + \sqrt{\frac{|b_0|^2 + |b_1|^2}{2}}|01> + \sqrt{\frac{|c_0|^2 + |c_1|^2}{2}}|10> + \sqrt{\frac{|d_0|^2 + |d_1|^2}{2}}|11> \quad (3)$$

This operation forces the probability amplitudes at the entrance of delay lines to be real valued and positive. Alternatively, we may use quantum summation (c-not gate).

3.1.4 Amplitudes to probabilities

The delays are fed with marginal probabilities of qubits, i.e. the probability of a qubit being measured as 0 if nothing is done on the other:

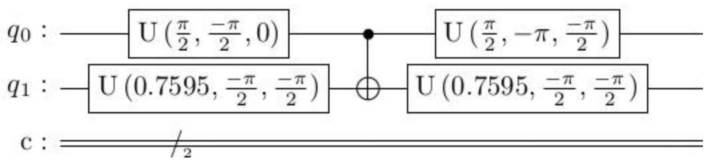

Figure 11: The 4×4 feedback matrix (2) as a quantum circuit.

$$p_{00} = \Pr(\mathsf{q}_0 = 0) = \sum_{b \in \{0,1\}} \Pr\big((\mathsf{q}_1, \mathsf{q}_0) = (b, 0)\big) = a_s^2 + c_s^2;$$

$$p_{10} = \Pr(\mathsf{q}_1 = 0) = \sum_{b \in \{0,1\}} \Pr\big((\mathsf{q}_1, \mathsf{q}_0) = (0, b)\big) = a_s^2 + b_s^2.$$

This is equivalent, for each qubit, to tracing out the other qubit from the density matrix and taking the top-left component.

3.1.5 Probabilities to amplitudes

The underdetermined system

$$a^2 + c^2 = p_{00}$$
$$a^2 + b^2 = p_{10},$$

constrained by the sum of squares of amplitudes being 1, is solved by

$$a^2 = min(p_{00}, p_{10})$$
$$b^2 = p_{10} - a^2$$
$$c^2 = p_{00} - a^2$$
$$d^2 = 1 + a^2 - p_{10} - p_{00}$$

3.1.6 Behavior

In response to a unit impulse, the QFDN of figure 10 produces an increasing density of echoes, as in a classical 2×2 FDN, but these appear to be superimposed to a fading-in square wave, with a periodicity that corresponds to the sum of the delay lengths (see figure 12).

If a measurement of one qubit is inserted after evolution of state (i.e., after matrix multiplication), the resulting sound is an irregularly-comb-filtered noise, as in the spectrogram of figure 13.

Another possibility is to repeatedly measure the state after matrix evolution, to transform frequencies of results to probabilities, and to derive a new state by taking the square root of such probabilities. The number of samples controls how noisy the output is, as referred to the case where probability amplitudes are used. The lower the number of samples, the higher the noisiness. With a very small number of samples (e.g., 8) the noise and its quantization become evident, as in figure 14, that should be compared with figure 12.

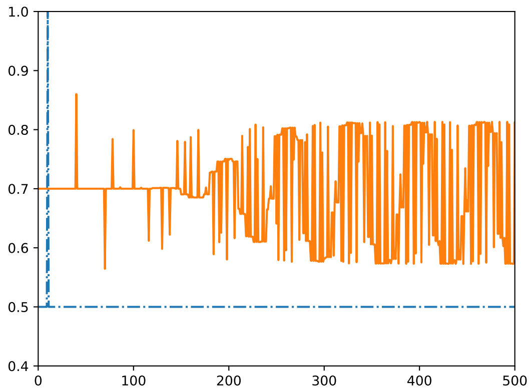

Figure 12: Response (continuous line) of the 2-qubit QFDN with feedback matrix (2), to a unit impulse. Time is measured in samples and delays are set equal to 29 and 37 samples.

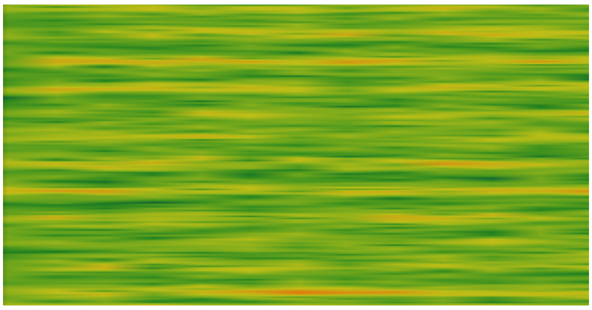

Figure 13: Spectrogram of the response of the 2-qubit QFDN with feedback matrix (2), to a unit impulse, under in-the-loop measurement of one qubit.

4 2^n delays and state-space realization

Similarly to how we did for the QCF in section 2.2, we can start from a QFDN of n qubits and propagate the 2^n probability amplitudes through 2^n delay lines. If all delays have the same length we are simply delaying a n-qubit quantum state, an operation that may be physically realizable. On the other hand, if the delays have different lengths, at the exit of the delay lines we are indeed composing a quantum state from probability amplitudes of differently-delayed states, and this makes sense only in simulation, where we know all the probability amplitudes.

In any case, however, a realization with 2^n delay lines can be transformed to state-space form and represented with a single large unitary matrix. How large it is depends on the lengths of the delay lines. For example, consider the case of $n = 2$ and 4 delay lines, and assume these lines have lengths $\{2, 2, 5, 7\}$. Considering the state as formed by the values at the exit of each delay unit, the state space will have size 16 and the overall unitary matrix will be 16×16. One possible state-space feedback matrix, embedding the circulant matrix $\mathbf{A}$ defined in (2), would be

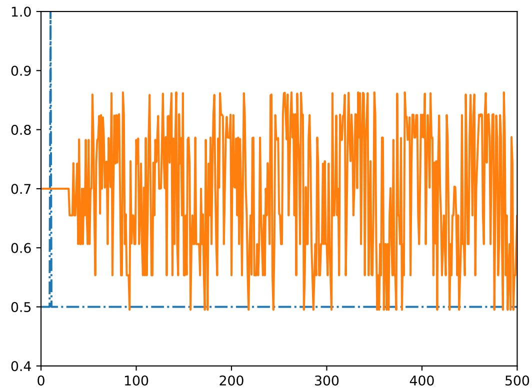

Figure 14: Response (continuous line) of the 2-qubit QFDN with feedback matrix (2), to a unit impulse. The state after matrix evolution is measured 8 times, and the results are converted to probabilities. Time is measured in samples and delays are set equal to 29 and 37 samples.

the unitary matrix

$$\mathbf{A}_{ss} =$$
$$\begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -.5 & .5 & .5 & .5 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & .5 & -.5 & .5 & .5 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & .5 & .5 & -.5 & .5 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & .5 & .5 & .5 & -.5 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix}$$

In state-space form, the QFDN can be represented by a feedback system with unit delays and quantum evolution on 4 qubits. Longer delays would require more qubits, and a correspondingly larger matrix that, however, may be realized as a quantum circuit.

5 Conclusion and further exploration

Starting from the observation that unitary matrices are among the energy-preserving matrices that can be used as scattering element in a feedback delay network, and driven by sonic curiosity, we conducted an exploration of possible quantum realizations of the FDN structure. We started with possible realizations of the 1-qubit recursive comb filter, with one or two delay lines. Propagating the probability amplitudes through different-length delay lines may not be physically feasible. However, the state-space realization of such structure may be actually seen as a quantum evolution of a larger number of qubits. These structures have been generalized to higher-order FDNs, such as those that are commonly used in artificial reverberation. In QFDNs, several qubits and delay lines are considered and with state-space realizations and long delays, the number of qubits and the size of the unitary operator may become high.

We have shown some examples of responses to impulse, damped sinusoid, or vocal signal, for different configurations, with or without measurement. We have seen how different the response may be from that of classical FDNs, although the behavior can be often interesting and controllable. With in-the-loop quantum measurement, the response is not deterministic and can produce non-repeating textures.

More structure design and experimentation need to be done, especially with state-space realizations that are physically realizable as quantum computations on several qubits. Exposing quantum operations, such as Bloch-sphere rotations, together with their auditory manifestations may give rise to novel controllable audio effects.

Acknowledgment

The presentation of this work has been supported by the project "Multiscale Analysis of Human and Artificial Trajectories: Models and Applications" funded by the MUR Progetti di Ricerca di Rilevante Interesse Nazionale (PRIN) Bando 2022 - grant 2022RB939W.

References

[1] Jean-Marc Jot and Antoine Chaigne. Digital delay networks for designing artificial reverberators. In *Audio Engineering Society Convention 90*. Audio Engineering Society, 1991.

[2] Ville Pulkki, Tapio Lokki, and Davide Rocchesso. Spatial effects. In Udo Zölzer, editor, *Digital Audio Effects*, pages 139–183. John Wiley and Sons, Ltd., Chichester Sussex, UK, 2011. Second edition.

[3] Davide Rocchesso. The ball within the box: A sound-processing metaphor. *Computer Music Journal*, 19(4):47–57, 1995.

[4] Sebastian J. Schlecht and Emanuël A. P. Habets. On lossless feedback delay networks. *IEEE Transactions on Signal Processing*, 65(6):1554–1564, 2017.

[5] Davide Rocchesso and Julius O. Smith. Circulant and elliptic feedback delay networks for artificial reverberation. *IEEE Transactions on Speech and Audio Processing*, 5(1):51–63, 1997.

[6] Paulo Vitor Itaboraí and Eduardo Reck Miranda. *Quantum Representations of Sound: From Mechanical Waves to Quantum Circuits*, pages 223–274. Springer International Publishing, Cham, 2022.

[7] Kevin Karplus and Alex Strong. Digital synthesis of plucked-string and drum timbres. *Computer Music Journal*, 7(2):43–55, 1983.

[8] M.R. Schroeder and B.F. Logan. "colorless" artificial reverberation. *IRE Transactions on Audio*, AU-9(6):209–214, 1961.

[9] James A. Moorer. About this reverberation business. *Computer Music Journal*, 3(2):13–28, 1979.

[10] Vesa Valimaki and Timo I. Laakso. Principles of fractional delay filters. In *2000 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.00CH37100)*, volume 6, pages 3870–3873 vol.6, 2000.

[11] Davide Rocchesso. Fractionally addressed delay lines. *IEEE Transactions on Speech and Audio Processing*, 8(6):717–727, 2000.

[12] Davide Rocchesso. Maximally diffusive yet efficient feedback delay networks for artificial reverberation. *IEEE Signal Processing Letters*, 4(9):252–255, 1997.