

DIFFERENTIABLE ATTENUATION FILTERS FOR FEEDBACK DELAY NETWORKS

Ilias Ibnyahya and Joshua D. Reiss,

Centre for Digital Music
Queen Mary University of London
London, UK

i.ibnyahya@qmul.ac.uk | joshua.reiss@qmul.ac.uk

ABSTRACT

We introduce a novel method for designing attenuation filters in digital audio reverberation systems based on Feedback Delay Networks (FDNs). Our approach uses Second Order Sections (SOS) of Infinite Impulse Response (IIR) filters arranged as parametric equalizers (PEQ), enabling fine control over frequency-dependent reverberation decay. Unlike traditional graphic equalizer designs, which require numerous filters per delay line, we propose a scalable solution where the number of filters can be adjusted. The frequency, gain, and quality factor (Q) parameters are shared parameters across delay lines and only the gain is adjusted based on delay length. This design not only reduces the number of optimization parameters, but also remains fully differentiable and compatible with gradient-based learning frameworks. Leveraging principles of analog filter design, our method allows for efficient and accurate filter fitting using supervised learning. Our method delivers a flexible and differentiable design, achieving state-of-the-art performance while significantly reducing computational cost.

1. INTRODUCTION

Feedback Delay Networks (FDNs) are a common technique in artificial reverberation, used to simulate the reflections found in real acoustic environments [1, 2]. They are especially used for modeling the late reflections of a reverberant field, where their scalable structure allows an effective balance between computational efficiency and acoustic accuracy. In typical FDNs, attenuation filters are employed to control the frequency-dependent decay, while the feedback matrix is designed to be both lossless and colorless [3, 4]. Each feedback path incorporates an absorption filter that gradually reduces energy over time. By using frequency-dependent filters, the reverberation decay can be precisely shaped across the spectrum, offering perceptually meaningful control over the reverberation time (T_{60}).

Attenuation filters, also referred to as absorption filters, simulate the loss of energy at different frequencies. While early methods relied on simple low-pass filters [5], more advanced designs now use graphic equalizers (GEQs) for finer spectral resolution [6, 7, 8]. Some recent work also explores deep learning approaches to directly learn IIR filter coefficients [9], but these methods typically lack a clear mechanism for uniformly controlling the decay time across all frequency bands. As a result, they may produce filters with inconsistent spectral decay characteristics, making them

less suitable for applications where precise and homogeneous reverberation control is required. Traditional Graphic Equalizers (GEQs) present limitations when integrated into Feedback Delay Networks (FDNs). Their typically limited attenuation range can result in poor matching within long delay lines. Furthermore, the inherent characteristics of the bell filters used in GEQs can introduce frequency response mismatches at both DC and Nyquist frequencies, as these filters must reach unity gain at the spectrum's edges [10]. The two-stage attenuation filter (TSAF) introduced in [8] combines a shelving filter with a 31-band GEQ to mitigate these issues and achieves strong performance with a low relative T_{60} error on real-world impulse responses [11]. Although a reduced version with 12 bands lowers computational cost, it can compromise accuracy. Furthermore, the granularity of these methods is tied to the number of bands per octave, making parametric equalizers a more computationally efficient alternative. Despite significant progress, achieving a fully differentiable and trainable FDN remains a challenge. Prior research has succeeded in optimizing gain values and feedback matrices [4, 12], but attenuation filters remain difficult to optimize in a differentiable manner [8, 9]. The challenge lies in the high dimensionality of the parameter space, as each delay line requires multiple filters and there is lack of a direct mapping between reverberation time and filter coefficients. Existing approaches often use fixed frequency and Q values, optimizing only the gain, and require a complex multi-stage processes. As such, they are not ideal for training in end-to-end differentiable frameworks. Another practical issue concerns the accuracy of digital filters across the full audio spectrum. Filter designs based on the bilinear transform, such as those in [13], tend to lose accuracy near the Nyquist frequency. Improved formulations like those of [14, 15] and the symmetric shelving design of [16] help maintain better precision, particularly in the high-frequency range.

To address these challenges, we propose a method that uses a variable amount of biquad filters per delay line, structured as a parametric equalizer with two shelving filters and bell filters. This setup delivers a magnitude response comparable to the more complex TSAF with 12 bands while reducing the number of filters by more than two-thirds. For an 8th-order FDN with a 12-bands PEQ, our method requires 96 biquads compared to 248 for a 31-band GEQ. Our system is fully differentiable and supports optimization through gradient descent, allowing all filter parameters to be trained jointly. This enables efficient and scalable filter fitting without compromising performance, which is particularly advantageous in resource-constrained environments such as embedded systems. An implementation of our approach is available on GitHub¹.

Copyright: © 2025 Ilias Ibnyahya et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, adaptation, and reproduction in any medium, provided the original author and source are credited.

¹https://github.com/ilias-audio/iir_match

2. METHOD

2.1. Parametric Equalizer Design

The parametric equalizer (PEQ) used in this work is scalable. It consists of three types of second-order filters: bell, low-shelf, and high-shelf. A typical configuration for digital reverberation is illustrated in Figure 1. The number of bands N and the delay line lengths m are variable. The first and last filters, $H_1(z)$ and $H_N(z)$, are implemented as low-shelf and high-shelf filters respectively, while the intermediate filters $H_2(z)$ through $H_{N-1}(z)$ are bell filters.




Figure 1: Recursive filter design with an N -band parametric equalizer.

Each filter is defined by three parameters: gain G , center frequency f_c , and quality factor Q . The gain is expressed in dB and converted to a linear amplitude factor A as follows:

$$A = \sqrt{10^{\frac{G}{20}}} = 10^{\frac{G}{40}} \quad (1)$$

We base our analog filter design on the general second-order prototype:

$$H(s) = g_{\infty} \frac{s^2 + \frac{\omega_z}{Q_z}s + \omega_z^2}{s^2 + \frac{\omega_p}{Q_p}s + \omega_p^2} \quad (2)$$

In this expression, g_{∞} is the high-frequency gain, while ω_z , Q_z and ω_p , Q_p define the zeros and poles respectively. Laplace domain transfer functions from [17] are used to define the response of Bell (H_B), Low-Shelf (H_{LS}) and High-Shelf (H_{HS}) filters:

$$H_B(s) = \frac{s^2 + \frac{A}{Q}s + 1}{s^2 + \frac{s}{AQ} + 1} \quad (3)$$

$$H_{LS}(s) = A \frac{s^2 + \frac{\sqrt{A}}{Q}s + A}{As^2 + \frac{\sqrt{A}}{Q}s + 1} \quad (4)$$

$$H_{HS}(s) = A \frac{As^2 + \frac{\sqrt{A}}{Q}s + 1}{s^2 + \frac{\sqrt{A}}{Q}s + A} \quad (5)$$

To evaluate the frequency response from the Laplace domain we substitute $s = j\frac{2\pi f}{f_c}$ and express the frequency response as a function of frequency f , with the center frequency f_c , linear gain A , and Q -factor Q as parameters:

$$|H_B(f)| = \sqrt{\frac{(1 - \frac{f}{f_c})^2 + (\frac{A}{Q})^2}{(1 - \frac{f}{f_c})^2 + (\frac{f}{AQ})^2}} \quad (6)$$

$$|H_{LS}(f)| = A \sqrt{\frac{(A - \frac{f}{f_c})^2 + (\frac{\sqrt{A}}{Q})^2}{(1 - A\frac{f}{f_c})^2 + (\frac{\sqrt{A}}{Q})^2}} \quad (7)$$

$$|H_{HS}(f)| = A \sqrt{\frac{(1 - A\frac{f}{f_c})^2 + (\frac{\sqrt{A}}{Q})^2}{(A - \frac{f}{f_c})^2 + (\frac{\sqrt{A}}{Q})^2}} \quad (8)$$

The magnitude response of the full PEQ is obtained by summing the logarithmic magnitude of each band:

$$20 \log_{10} |H_{PEQ}(f)| = \sum_{i=1}^N 20 \log_{10} |H_i(f)| \quad (9)$$

Each filter band thus requires optimization over three parameters. The analog prototype is converted to a digital filter using the method in [14].

Following the approach of Prawda et al. [3], we model the attenuation response to scale with the delay line length in samples. The target magnitude for a given delay line k of length m_k and sample rate f_s becomes:

$$20 \log_{10} |H_{PEQ_k}(f)| = \frac{-60 m_k}{T_{60}(f) f_s} \quad (10)$$

This relation ensures that all filters can share a common set of parameters while scaling the gain according to delay line length. We define the slope as:

$$\gamma(f) = \frac{-60}{T_{60}(f)} \quad (11)$$

and the corresponding gain per delay line is:

$$G(f) = \frac{\gamma(f) m_k}{f_s} \quad (12)$$

Using this formulation, the PEQ can be shared across all delay lines with only the gain scaling changing, significantly reducing the number of parameters.

2.2. Parameter Optimization

To optimize the filter parameters, we frame the task as supervised regression. The objective is to match the frequency-dependent target magnitude derived from a desired $T_{60}(f)$. The loss function is the mean squared error (MSE) between the target and the PEQ response:

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N \left(\frac{-60 m_i}{T_{60}(f) f_s} - 20 \log_{10} |H_{PEQ}(f)| \right)^2 \quad (13)$$

We implement the full model in PyTorch, where each filter band is differentiable and parameterized by frequency, gain, and Q -factor. Optimization is performed with the Adam optimizer [18], using an initial learning rate of 0.1 and 10,000 iterations. These values are empirical but provide enough iterations for complex cases with more than 12 bands.

At each step, gradients are computed with respect to all parameters via backpropagation, allowing efficient training of the full PEQ. The use of shared parameters across delay lines ensures scalability and enables the application of our method even in embedded systems.

Finally, we validate the method using a dataset of 1000 RIRs [11], following the evaluation protocol from [8], and compare it against state-of-the-art approaches.




Figure 2: Reverberation time relative error of three different PEQ designs applied to a 1000 RIR reverberation time measurements dataset with a distribution of different delay lengths.




Figure 3: Magnitude response accuracy for different PEQ designs based on the dataset median reverberation time at each frequency.

3. RESULTS

3.1. Reverberation Time Evaluation

We evaluated the precision of our parametric equalizer designs against a dataset of measured room impulse responses (RIR) in a variable acoustics environment [19]. The reverberation times are estimated in third-octave frequency bands over 31 bands. A subset of 1000 RIRs is used, following a comparable method to the one exposed in [8]. We linearly interpolate the 31 bands of RIR measurements to 512 logarithmically spaced points between DC and Nyquist. Our approach is tested by randomly generating delay line lengths between 0.01 and 0.3 seconds to emulate the behavior of different types of rooms and halls. This strategy also allows for evaluation against a wide range of attenuation filter combinations, where, for example, an increase in delay corresponds to greater attenuation. Each test case involved matching the frequency dependent reverberation time ($T_{60}(f)$) with the optimized magnitude response of the filter converted back to a reverberation time ($\hat{T}_{60}(f)$).

$$T_{60, \text{error}} = \frac{T_{60}(f) - \hat{T}_{60}(f)}{T_{60}(f)} \times 100\%. \quad (9)$$

We calculate the error $T_{60, \text{error}}$ for each point of the interpolated dataset and for each of the 1000 reverberation time estimations. This effectively means we evaluate the error against 512,000 data points. We then use this relative error to determine the error distribution (Figure 2), observing that high levels of inaccuracy are avoided. For instance, the smallest equalizer (EQ) yields an inaccuracy within $\pm 25\%$. 12 bands PEQ shows a narrow distribution of the error centered around zero. Adding more bands makes the resulting frequency reverberation time attenuation more accurate.

In Figure 3, even with as few as 4 bands, the model captures the general trend of the target magnitude curve. Increasing the number of bands to 8 and 12 further reduces the spread of the error, bringing the performance close to that of the state-of-the-art Two-Stage Attenuation Filter (TSAF) with 31 bands. This highlights the effectiveness of our differentiable approach in capturing complex reverberation behavior with far fewer filters.

3.2. Attenuation Filter Magnitude Accuracy

To assess the fidelity of the frequency response, we measured the deviation between the optimized equalizer response and the median target T_{60} curve across the dataset. We used a delay line of 100 milliseconds and a sampling rate of 48kHz. Figure 3 displays the achieved magnitude response for various PEQ configurations.

The results show that even low-order designs (e.g., 4 or 8 bands) provide a smooth and perceptually relevant approximation of the target response. The 12-band PEQ closely tracks the target curve, exhibiting only minor deviations, and is nearly indistinguishable from the target.

Type	$MSE(x, \hat{x})$	$MAE(x, \hat{x})$	OP	P
PEQ_4	4.8×10^{-2}	6.3×10^{-1}	36	12
PEQ_8	7.3×10^{-3}	4.1×10^{-1}	72	24
PEQ_{12}	1.7×10^{-3}	2.2×10^{-1}	108	36
$TSAF_{31}$	1.9×10^{-3}	1.2×10^{-1}	284	33

Table 1: Magnitude Mean-Squared Error (MSE) and Maximum Absolute Error (MAE) in dB between the target response (x) and the predicted response (\hat{x}) using N -band PEQs and the 31-band TSAF. OP denotes the number of arithmetic operations, and P the number of trainable parameters.

Quantitative metrics are presented in Table 1, comparing the Mean Squared Error (MSE), Maximum Absolute Error (MAE), and computational cost (additions and multiplications, noted as OP) for each method using the same delay length and sampling rate. Additionally, P represents the number of trainable parameters, providing insight into the complexity and memory footprint of each configuration. Our proposed 12-band PEQ achieves a similar MSE to the TSAF while using only 38% of the operations and maintaining a comparable number of parameters (36 for our approach and 33 for $TSAF_{31}$), indicating a favorable balance between performance, computational efficiency, and model size for the proposed architecture. These findings confirm that our differentiable PEQ design offers a balance between performance and complexity, making it suitable for real-time applications.

4. CONCLUSIONS

We presented a simple and efficient method for designing attenuation filters in Feedback Delay Networks using a differentiable parametric equalizer framework. Our approach uses fewer filters than traditional methods while maintaining comparable accuracy and stability, making it well-suited for real-time and embedded applications. By optimizing the filters with gradient descent and working on continuous target curves, we avoid the need for complex neural networks.

Future work should explore integrating this method into a complete FDN architecture and evaluate its effectiveness using biquad coefficients. This would help assess its performance in realistic audio processing pipelines and confirm its suitability for end-to-end optimization tasks.

5. REFERENCES

- [1] Vesa Välimäki, Julian D. Parker, Lauri Savioja, Julius O. Smith, and Jonathan S. Abel, “Fifty years of artificial reverberation,” *IEEE Transactions on Audio, Speech and Language Processing*, vol. 20, no. 5, pp. 1421–1448, July 2012.
- [2] J.-M. Jot, “An analysis/synthesis approach to real-time artificial reverberation,” in *[Proceedings] ICASSP-92: 1992 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Mar. 1992, vol. 2, pp. 221–224 vol.2.
- [3] Karolina Prawda, Vesa Välimäki, and Sebastian J. Schlecht, “Improved reverberation time control for feedback delay networks,” in *Proceedings of the 22nd International Conference on Digital Audio Effects (DAFx-19)*, Birmingham, UK, Sept. 2019.
- [4] Gloria Dal Santo, Karolina Prawda, Sebastian Schlecht, and Vesa Välimäki, “Differentiable Feedback Delay Network For Colorless Reverberation,” in *Proceedings of the 26th International Conference on Digital Audio Effects (DAFx23)*, Copenhagen, Denmark, Sept. 2023, pp. 244–251, Aalborg University.
- [5] James A. Moorer, “About This Reverberation Business,” *Computer Music Journal*, vol. 3, no. 2, pp. 13–28, 1979.
- [6] Jean Marc Jot, “Proportional parametric equalizers - Application to digital reverberation and environmental audio processing,” in *139th Audio Engineering Society International Convention*, New York, USA, Oct. 2015, AES.
- [7] Vesa Välimäki and Juho Liski, “Accurate cascade graphic equalizer,” *IEEE Signal Processing Letters*, vol. 24, no. 2, pp. 176–180, 2017.
- [8] Vesa Välimäki, Karolina Prawda, and Sebastian Schlecht, “Two-Stage Attenuation Filter for Artificial Reverberation,” *Signal Processing Letters, IEEE*, vol. 31, pp. 391–395, Jan. 2024.
- [9] Sungho Lee, Hyeong-Seok Choi, and Kyogu Lee, “Differentiable Artificial Reverberation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 2541–2556, 2022.
- [10] Juho Liski and Vesa Välimäki, “The Quest for the Best Graphic Equalizer,” in *Proceedings of the 20th International Conference on Digital Audio Effects (DAFx-17)*, Edinburgh, UK, Sept. 2017.
- [11] Karolina Prawda, Sebastian J. Schlecht, and Vesa Välimäki, “Calibrating the Sabine and Eyring formulas,” *The Journal of the Acoustical Society of America*, vol. 152, no. 2, pp. 1158–1169, Aug. 2022.
- [12] Gloria Dal Santo, Benoit Alary, Karolina Prawda, Sebastian Schlecht, and Vesa Välimäki, “RIR2FDN: An improved room impulse response analysis and synthesis,” in *Proceedings of the 27th International Conference on Digital Audio Effects (DAFx24)*, University of Surrey, Sept. 2024, pp. 230–237.
- [13] Shahan Nercessian, “Neural Parametric Equalizer Matching Using Differentiable Biquads,” *Proceedings of the 23rd International Conference on Digital Audio Effects (DAFx2020)*, , no. September, pp. 265–272, 2020, Number: September.
- [14] Juho Liski, Jussi Rämö, and Vesa Välimäki, “Graphic Equalizer Design with Symmetric Biquad Filters,” in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2019, pp. 55–59.
- [15] Flynn John and Joshua D. Reiss, “Improving the frequency response magnitude and phase of analogue-matched digital filters,” in *144th Audio Engineering Society International Convention*, May 2018, AES.
- [16] Vesa Välimäki and Joshua D. Reiss, “All about audio equalization: Solutions and frontiers,” *Applied Sciences*, vol. 6, no. 5, 2016.
- [17] Robert Bristow-Johnson, “Audio-EQ-Cookbook,” 2021, available at “<https://www.w3.org/TR/audio-eq-cookbook/>” accessed June 08, 2025.
- [18] Diederik P. Kingma and Jimmy Lei Ba, “Adam: A Method for Stochastic Optimization,” in *3rd International Conference for Learning Representations*, San Diego, 2015.
- [19] Karolina Prawda, Sebastian J. Schlecht, and Vesa Välimäki, “Dataset of impulse responses from variable acoustics room Arni at Aalto Acoustic Labs,” Aug. 2022.