Parametric FIR Design of Propagation Loss Filters in Digital Waveguide String Models

Published in:
IEEE Signal Processing Letters

Document Version:
Peer reviewed version

Queen's University Belfast - Research Portal:
Link to publication record in Queen's University Belfast Research Portal

Parametric FIR Design of Propagation Loss Filters in Digital Waveguide String Models

Maarten van Walstijn

Abstract—One of the attractive features of sound synthesis by physical modeling is the potential to build acoustic-sounding digital instruments that offer more flexibility and different options in its design and control than their real-life counterparts. In order to develop such virtual-acoustic instruments, the models they are based on need to be fully parametric, i.e. all coefficients employed in the model are functions of physical parameters that are controlled either on-line or at the (off-line) design stage. In this letter we show how propagation losses can be parametrically incorporated in digital waveguide string models with the use of zero-phase FIR filters. Starting from the simplest possible design in the form of a three-tap FIR filter, a higher-order FIR strategy is presented and discussed within the perspective of string sound synthesis with digital waveguide models.

EDICS: AEA-AUEA Audio and Electroacoustics, DSP-FILT Filter design, analysis and implementation

I. INTRODUCTION

DIGITAL waveguides have been used for efficient simulation and sound synthesis of 1-D musical resonators, such as string and pipes, for more than 20 years [1]. A basic digital waveguide (DW) structure consists of a pair of delay-lines that directly simulate the traveling wave solutions of the 1-D wave equation. For realistic sound synthesis of string instruments, various physical phenomena, such as propagation losses, end reflections, and where appropriate, stiffness, or even non-linear phenomena such as tension-modulation [2] must be taken into account. The focus of this letter is on the incorporation of propagation losses, which in principle can be done by inserting loss filters in the basic DW structure (see Figure 1). The concept of using loop filters to account for losses in string simulations originated in studies on the Karplus-Strong algorithm [3], [1], a well-known forerunner of digital waveguides.

A common approach is to lump the propagation losses together with those occurring at the bridge, and modeling all losses with a single loop filter [4]. Due to the typical complexity of the bridge reflection function, the overall loss filter is then usually designed via optimisation, often fitting to experimental data (see, e.g. [5], [6]). This approach works well when the objective is to simulate a 'fixed instrument' with prescribed specifications, but does not allow easy adjustment of the characteristics of either the string or the bridge/body. Such separate control is of interest when the objective is not to approximate a prescribed string response, but instead to

M. van Walstijn is with the Sonic Arts Research Centre (SARC) and the School of Electronics, Electrical Engineering and Computer Science, Queen's University, Belfast, United Kingdom BT7 1NN e-mail: (m.vanwalstijn@qub.ac.uk)

more widely explore the sonic potential of a virtual-acoustic string instrument, either in real-time or off-line; this includes venturing into domains where the string losses are artificially high or low, thus making higher demands regarding parameter ranges. Consequently, the propagation loss filter must remain separate and physically parametric, i.e. its coefficients are defined and calculated as closed-form functions of one or more physical parameters. This allows 'tuning' the string losses, similar to adjusting the string tension with tunable fractional delay filters [7], or the string stiffness with tunable allpass filters [8]. Commonly used tunable loss-filter include first-order FIR [9] and IIR [10], [9] designs, the coefficients of which can be linked to physical parameters [9]. A disdavantage of these designs is that the phase delay response varies with the loss parameters, causing the fundamental frequency of the string to change with one of the loss parameters. This problem can be avoided by using a linear-phase FIR design, such as the second-order filters proposed in [11], [1]. However no direct link to the physical loss parameters is given in these, and in addition they have a limit on the spectral roll-off, and therefore on the amount of frequency-dependent loss that can be achieved.

The aim of this letter is to design tunable propagation loss filters, where the main criterion is low-cost filters with simple coefficient formulas. Further properties assumed to be useful are direct control over accuracy via filter order, a perceptually sensible dependence of accuracy on frequency, filters that do not affect the frequencies of the partials (so that this can be controlled separately), and the ability to approximate any specific instance of the theoretical target response.

II. LOSS FILTER TARGET RESPONSE

Propagation losses in strings can generally be attributed to two physical phenomena, namely the loading by the surrounding air and internal friction [12]. These effects can be modelled loosely by adding two loss terms to the transversal wave equation, that for a flexible string then takes the form [13]:

$$\frac{\partial^2 u}{\partial t^2} = c^2 \frac{\partial^2 u}{\partial x^2} - 2b_1 \frac{\partial u}{\partial t} + 2b_2 \frac{\partial^3 u}{\partial x^2 \partial t} \qquad (1)$$

where u is the transversal displacement, c is the wave speed, and where the constants b_1 and b_2 are the loss parameters. It is straightforward to show that displacement waves governed by (1) decay at a rate $\alpha = b_1 + b_2 k^2$, where k is the wave number, that relates to the frequency by $\omega^2 = k^2 c^2 - \alpha^2$ [13]. The same holds for velocity waves, that are commonly employed in DW models. Since for real strings, $\alpha \ll k^2 c^2$, it is reasonable to

use instead the wave number for ideal strings ($k = \omega/c$), so that for propagation over distance d and time $\tau = d/c$, the target response of the digital loss filter can be defined as

$$H_L(\omega) = ge^{-b_2 \, c^{-2} \tau \omega^2}, \tag{2}$$

where $g = e^{-b_1 \tau}$ is a simple gain factor. Hence a suitable target response for a digital loss filter that approximates the frequency-dependent part of H_L, i.e. $H_L(\omega)/g$, can be defined as

$$H(\Omega) = e^{-\beta \Omega^2}, \tag{3}$$

where $\beta = (b_2 \tau)/(c^2 T^2)$ is a convenient design parameter and $\Omega = \omega T$ denotes digital frequency for sampling period T. The Taylor expansion of (3) about $\Omega = 0$, on which all FIR designs in this letter are based, is

$$H(\Omega) = 1 + \sum_{i=1}^{\infty} \frac{\beta^i}{i!}(-\Omega^2)^i \tag{4}$$

In order to assess the accuracy of any digital filter designed to match the target response, one may compare the digital filter frequency response $\hat{H}(\Omega)$ to $H(z)$. Alternatively, the filter's effective β value can be computed as a function of frequency as follows

$$\hat{\beta}(\Omega) = -\frac{\ln \hat{H}(\Omega)}{\Omega^2} \tag{5}$$

and compared to the theoretical target value β. The latter comparison is used throughout the letter since it allows a more clear visual inspection.

III. THREE-TAP FIR DESIGN

The target response $H(\Omega)$ is zero-phase, so it makes sense to design a zero-phase digital loss filter [1]; this can be achieved using an N-tap FIR with symmetrical coefficients, and 'stealing' $M = (N-1)/2$ delays from the delay-line. Regarding the DW structure in Figure 1, the lowest number of delays in each of the delay-lines is one, which corresponds to having only one modal frequency (i.e. the fundamental) below Nyquist. In that case, propagation losses may be modelled efficiently using only a gain factor g. For two or more modal frequencies below Nyquist, the structure will have at least two delays in each of the delay-lines. Hence a three-tap, second-order zero-phase FIR filter, that requires stealing only a single delay, is always realisable where needed, and therefore a logical first design choice. The transfer function of this filter can be parameterized suitably as

$$\hat{H}(z) = \theta z + (1 - 2\theta) + \theta z^{-1}, \tag{6}$$

where θ is a free parameter that controls the spectral roll-off. The frequency response is

$$\hat{H}(\Omega) = 1 + 2\theta(\cos \Omega - 1), \tag{7}$$

and its Taylor approximation about $\Omega = 0$ is

$$\hat{H}(\Omega) = 1 + 2\theta_m \sum_{i=1}^{\infty} \frac{(-\Omega^2)^i}{(2i)!} \tag{8}$$

Comparing (8) to (4) for $i = 1$ immediately reveals that a second-order accuracy approximation is possible by setting

$\theta = \beta$. Plotting $\hat{H}(\Omega)$ and $\hat{\beta}(\Omega)$ against normalized frequency ($fT = \Omega/(2\pi)$) for a range of β values (see Figure 2) confirms that accuracy is high at lower frequencies and decreasing with frequency, but also reveals that the three-tap FIR design works well only for $\beta \leq 0.25$; for higher values, the effective damping is too small, leading to high-frequency modes decaying far too slowly. Such high-frequency ringing can become perceptually significant and would present an unpleasant model artefact, and should therefore be avoided. Hence the three-tap design has a distinct limit on β.

IV. N-TAP FIR DESIGN

For audio-rate simulations where τ is considerably larger than $2T$, which is the case for almost all real musical strings, higher-order accuracy can be achieved by using a larger FIR filter, the generalized transfer function of which takes the form

$$\hat{H}(z) = 1 + \theta_m \sum_{m=1}^{M} \left[z^m - 2 + z^{-m} \right] \tag{9}$$

where we now have M free parameters θ_m. The corresponding FIR frequency response is

$$\hat{H}(\Omega) = 1 + 2 \sum_{m=1}^{M} \theta_m \left[\cos m\Omega - 1 \right] \tag{10}$$

and its Taylor expansion can be written

$$\hat{H}(\Omega) = 1 + \sum_{i=1}^{\infty} \frac{-a_i}{i!}(-\Omega^2)^i \tag{11}$$

where

$$a_i = \frac{2i!}{(2i)!} \sum_{m=1}^{M} \theta_m m^{2i} \tag{12}$$

Comparing (11) to (4) reveals that accuracy of order $2M$ is achieved when $a_i = \beta^i$ for $i = 1, 2, 3..M$. This amounts to solving a linear system of M equations, which in principle can be directly applied to calculate the free parameters θ_m (see accompanying Matlab file NtapFIR.m). In practice however it is convenient to pre-calculate closed-form expressions for specified orders. Table I lists these formulae for $M = 1, 2, 3$, and Figure 3 plots the resulting effective β value responses for $M = 2, 3$.

As could be expected, using a higher filter order increases accuracy, and also stretches the valid β range somewhat. However, for higher values of β, the effective damping at high frequencies is still heavily under-approximated, and further plots (not shown here) showed that the damping can even become negative, which would lead to instability. An exact limit on β is not as easy to define now, but in general it can be said that, for a filter using $N = 2M + 1$ taps, it lies considerably lower than M times the limit of 0.25 for a basic three-tap filter.

V. SERIES FIR DESIGN

In attempting to find a better way of overcoming the limit on β, it is worthwhile realizing that a steeper spectral roll-off

can be achieved simply by cascading several identical second-order FIR sections of the form of (6). The transfer function for L FIR sections in series then thus takes the form

$$\hat{H}(z) = \left[1 + \theta(z - 2 + z^{-1})\right]^L \qquad (13)$$

Given that there is now only a single free parameter (θ), the maximum accuracy that can be achieved is second-order. Applying the Taylor series expansion again and truncating all terms higher than second-order gives

$$\hat{H}(\Omega) \approx (1 - \theta\Omega^2)^L \approx 1 - L\theta\Omega^2, \qquad (14)$$

Therefore, in order to approximate (4), we now have to set $\theta = \beta/L$. Since the limit for a single section is still $\theta \leq 0.25$, the maximum value of β for which the filter behaves well at higher frequencies is now $L/4$, i.e. L times higher than for a single tree-tap FIR section. It follows that for a given value of β, the required filter order can be determined by setting $L = \lceil 4\beta \rceil$. Figure 4 shows the effective β value of a two-section series FIR design for a number of β values in the valid β range.

VI. Series of N-tap Sections

The fact that one may extend the valid β range simply by cascading identical FIR sections could of course also have been seen directly from nature of the target response function in (4). Hence this idea holds generally, and may thus also be applied to N-tap sections. That is, we may also cascade L instances of a N-tap FIR section, each of which approximates β in order to approximate a target value $L\beta$. This combinatorial approach allows fine-tuning the balance between 'bandwidth of accuracy' and 'maximum β', and generalizes the FIR design strategy proposed in this letter. The accompanying Matlab file `plotPPLFIR.m` can be used to generate plots for $\hat{H}(z)$ or $\hat{\beta}(z)$ for any choice of β, M, L.

VII. Concluding Remarks and Future Directions

In this letter we have discussed methods for designing physically parametric loss filters for use in digital waveguide string models. A zero-phase FIR approach was proposed, with the FIR coefficients derived as functions of the loss parameters by enforcing equivalence between the FIR frequency response and the theoretical target response up to a certain Taylor expansion order. As could be expected from this approach, the resulting filters are exact at $\omega = 0$ and the accuracy generally decreases with frequency. From a perceptual perspective, this is a useful attribute, given that in the digital waveguide model output signal the lower frequencies are generally more prominent and slower decaying than high frequencies. However, large under-approximations of the effective damping should be avoided as it would result in noticeable artefacts; this is what sets a limit on the size of the design parameter β.

In many applications, β is relatively small (typically below 0.1 for a 44.1kHz sample rate) and, given that the perceptual tolerance for decay parameters is typically high [14], high accuracy is often not required. In that case, a single three-tap FIR design is likely to suffice. The proposed design differs from previous three-tap FIR designs [11], [1] in that

the coefficients are calculated from the parameters of the underlying partial differential equation.

When higher accuracy is needed, the N-tap FIR design can be used. This may occur, for example, in inverse filtering applications (e.g. [4]) or quantitative comparisons of DW modeling with other physical modeling paradigms (e.g. [15]). It is worthwhile noting though that the target response is itself based on a fairly simplified, smooth model of the propagation losses, hence it remains to be verified to what extent a strong match with this parameterization means something toward synthesis realism or estimation accuracy.

For applications in which it is envisaged that β is a time-varying parameter, the series FIR design is attractive, since a large valid β range can be achieved with a lower filter order than with the N-tap design. In addition, this design uses only a single parameter that can be updated at very little cost. Hence the series FIR design is greatly suited to real-time implementations in which damping is one of the control parameters. Accuracy can again be improved by using an N-tap FIR filter as the nominal section to be cascaded.

An interesting option to explore in future work is to extend the loss parameterisation itself, i.e. adding more mixed-term derivatives to equation (1), and investigate whether this allows an improved fit to, for example, measured data on propagation losses in strings, or to the more complex case of wall losses in pipes. If successful, it seems possible then to directly apply the principles of the FIR design approach presented in this letter to these cases.

References

[1] J. O. Smith, *Physical Audio Signal Processing*. Stanford University, http://ccrma.stanford.edu/ jos/pasp/, accessed May 2010, online book.

[2] B. Bank, "Energy-based synthesis of tension modulation in strings," in *12th International Conference on Digital Audio Effects (DAFX09), Como, Italy*, 2009, pp. 365 –372.

[3] D. A. Jaffe and J. O. Smith, "Extensions of the Karplus-Strong plucked string algorithm," *Computer Music Journal*, vol. 7, no. 2, pp. 56–69, 1983.

[4] V. Välimäki, J. Huopaniemi, Karjalainen, and Z. M., Janosy, "Physical modeling of plucked string instruments with application to real-time sound synthesis," *Journal of the Audio Engieering Society*, vol. 44, no. 5, pp. 331 – 353, 1996.

[5] B. Bank and V. Välimäki, "Robust loss filter design for digital waveguide synthesis of string tones," *Signal Processing Letters, IEEE*, vol. 10, no. 1, pp. 18 – 20, jan. 2003.

[6] H.-M. Lehtonen, Rauhala, J., and V. Välimäki, "Sparse multi-stage loss filter design for waveguide piano synthesis," in *Applications of Signal Processing to Audio and Acoustics, 2005. IEEE Workshop on*, 16-16 2005, pp. 331 –334.

[7] V. Välimäki and T. Laakso, "Principles of fractional delay filters," in *Acoustics, Speech, and Signal Processing, 2000. ICASSP '00. Proceedings. 2000 IEEE International Conference on*, vol. 6, 2000, pp. 3870 –3873 vol.6.

[8] J. Rauhala and V. Välimäki, "Tunable dispersion filter design for piano synthesis," *Signal Processing Letters, IEEE*, vol. 13, no. 5, pp. 253 – 256, may 2006.

[9] B. Bank, "Physics-based sound synthesis of the piano," Master's thesis, Budapest Univ. Technol. and Economics, Budapest, Hungary, 2000, report 54, Helsinki Univ. of Technology, Lab. of Acoustics and Audio Signal Processing.

[10] V. Välimäki and T. Tolonen, "Development and calibration of a guitar synthesizer," *J. Audio Eng. Soc.*, vol. 46, no. 9, pp. 766 – 778, 1998.

[11] C. Cuzzucoli and V. Lombardo, "A physical model of the classical guitar, including the players touch," *Computer Music Journal*, vol. 23, no. 2, pp. 52 – 69, 1999.

[12] N. Fletcher and T. Rossing, *The Physics of Musical Instruments*. Springer, 1998, second edition.

[13] J. Bensa, S. Bilbao, R. Kronland-Martinet, and J. O. S. III, "The simulation of piano string vibration: From physical models to finite difference schemes and digital waveguides," *Journal of the Acoustical Society of America*, vol. 114, no. 2, pp. 1095–1107, August 2003.

[14] H. Järveläinen and V. Välimäki, "Perceptual tolerances for decay parameters in plucked string synthesis," *J. Audio Eng. Soc.*, vol. 49, no. 11, pp. 1049 – 1059, 2001.

[15] C. Erkut and M. Karjalainen, "Finite difference method vs. digital waveguide method in string instrument modeling and synthesis," in *International Symposium on Musical Acoustics, Mexico City, Mexico*, 2002.

TABLES AND FIGURES

TABLE I
PARAMETER FORMULAE FOR N-TAP FIR DESIGN.

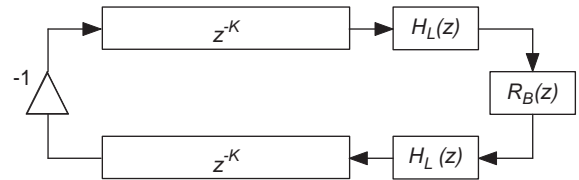| M	1	2	3
θ_1	β	$\frac{4}{3}\beta - 2\beta^2$	$\frac{3}{2}\beta - \frac{13}{4}\beta^2 + \frac{5}{2}\beta^3$
θ_2		$\frac{1}{12}\beta - \frac{1}{2}\beta^2$	$-\frac{3}{2}\beta + \beta^2 - \beta^3$
θ_3			$\frac{1}{90}\beta - \frac{1}{12}\beta^2 + \frac{1}{6}\beta^3$

Fig. 1. A digital waveguide model of a lossy string. Wave propagation is simulated with a fractional delay-line z^{-K} in series with a loss filter $H_L(z)$. The minus one reflection at the left end models a fixed boundary, and the filter $R_B(z)$ models the reflection from the bridge.
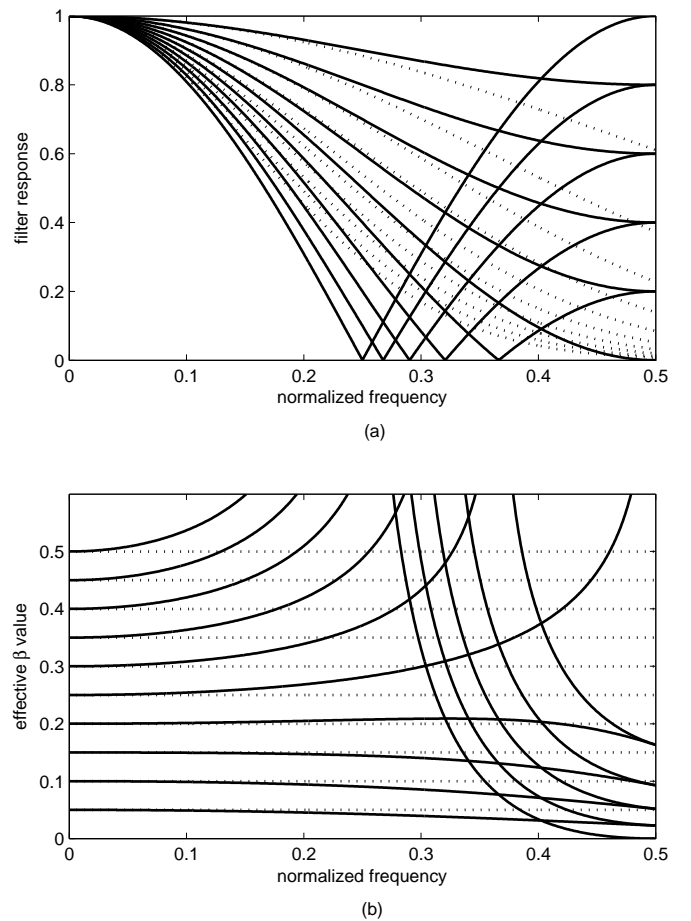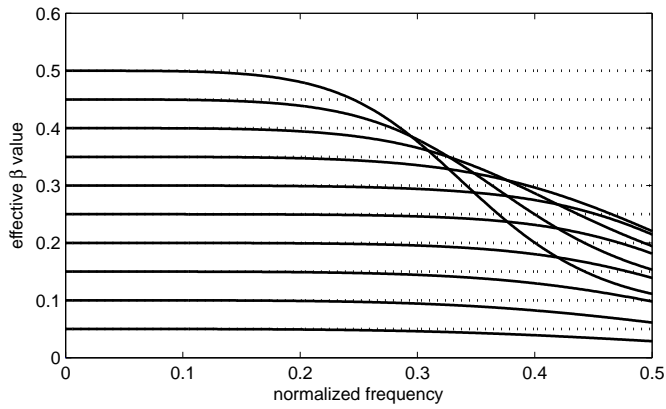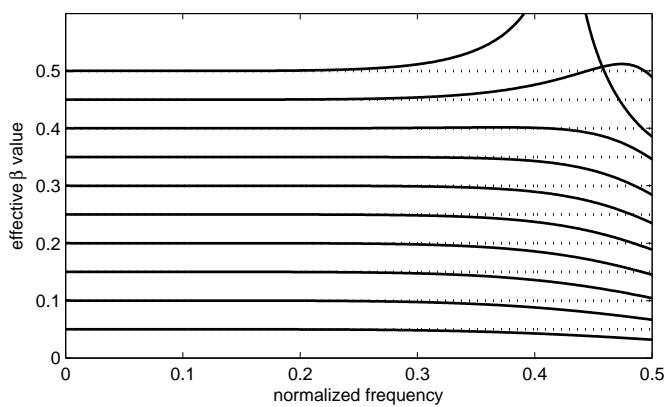

(a)

(b)

Fig. 2. The frequency response (a) and effective β value (b) of the three-tap FIR loss filter (solid) compared to the target response (dotted), for $\beta = 0.05, 0.10, 0.15...0.5$.

Fig. 3. The effective β value of (a) the five-tap ($M = 2$) FIR design, and (b) the seven-tap ($M = 3$) FIR design compared to the target response (dotted), for $\beta = 0.05, 0.10, 0.15...0.5$.
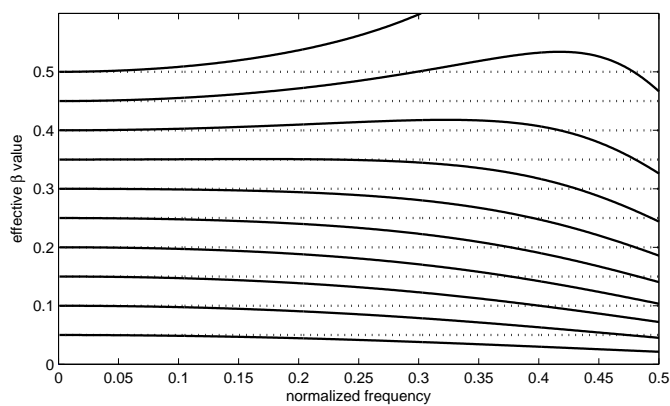

Fig. 4. The effective β value of an $L = 2$ series FIR design (solid) compared to the target values (dotted), for $\beta = 0.05, 0.10, 0.15...0.5$.