Audio Engineering Society

Convention Paper

Presented at the 132nd Convention
2012 April 26–29 Budapest, Hungary

Magnitude-priority filter design for audio applications

Balázs Bank[1]

[1] *Budapest University of Technology and Economics, Department of Measurement and Information Systems, Hungary*

Correspondence should be addressed to Balázs Bank (bank@mit.bme.hu)

ABSTRACT

In audio, often specialized filter design methods are used that take into account the logarithmic frequency resolution of hearing. A notable side-effect of these quasi-logarithmic frequency design methods is a high-frequency attenuation for non-minimumphase targets due to the frequency-dependent windowing effect of the filter design. This paper presents two approaches for the correction of this high-frequency attenuation, based either on the iterative update of the magnitude, or the iterative update of the phase of the target specification. As a result, the filter follows both magnitude and phase in those frequency regions where it can, while where this is not possible, it focuses on the magnitude. Thus, the new method combines the advantages of traditional complex and magnitude-only filter designs. The algorithms are demonstrated by parallel filter designs, but since the method does not make any assumption on the filter design algorithm used in the iteration, it is equally applicable to other techniques, like standard FIR, IIR, warped FIR, warped IIR, or Kautz filters.

1. INTRODUCTION

Most FIR and IIR filter design methods minimize the error between the filter frequency response and the complex target response, which we will call "straightforward" or "complex" methods in the paper. This is implemented directly in frequency-domain methods, but time-domain filter design methods also result in matching the complex target frequency response, since minimizing the mean squared error between the target and filter impulse responses is equivalent to complex frequency response matching due to Parseval's theorem. This means that when matching is perfect, complex methods result in a filter that follows the magnitude and phase of the target precisely. However, if perfect matching cannot be achieved, e.g., due to low filter order, this will result in both amplitude and phase errors.

In some applications the magnitude response of the filter is much more critical than the phase response: in this case "magnitude-only" design techniques can be used. The usual way of implementing this is making the target response minimum-phase prior to complex filter design, that is, altering the target phase in such a way that is easy to follow for the design algorithm. As a result, the magnitude matching will be more accurate compared to complex design of the same order, but the filter phase might be very different from the phase of the original target.

This paper presents a new alternative that combines the advantages of complex and magnitude-only filter designs. The basic idea of "magnitude-priority" design is that it tries to model the complex transfer function whenever possible, but in those frequency regions where this cannot be done, it concentrates on the magnitude. As a result, the magnitude matching will be just as accurate as with magnitude-only methods, while the phase response will get closer to the target.

Although the magnitude-priority filter design may present some advantages over the magnitude-only methods in many fields of signal processing, it is particularly well suited for audio applications because of the properties of quasi-logarithmic filter design methods discussed in the next section.

2. QUASI-LOGARITHMIC FILTER DESIGN METHODS

Efficient audio filtering and equalization requires that filter design takes into account the quasi-logarithmic frequency resolution of the hearing, as opposed to the linear frequency resolution of traditional FIR and IIR filter design techniques. Therefore, special filter design methodologies have been developed, including warped [1], Kautz [2], and parallel [3] filters that allow filter design at a logarithmic-like frequency resolution. It has been shown that these quasi-logarithmic filter design methods result in a filter response that is similar to the frequency-dependent windowing (or, equivalently, transfer function smoothing) of transfer functions [4, 5]. The effective length of the window is inversely proportional to the frequency resolution Δf. Thus, a logarithmic frequency resolution, where Δf is proportional to the frequency f, results in a window length that decreases as a function of frequency $1/f$.

A notable side-effect of frequency-dependent windowing is that the high-frequency response is attenuated for non-minimumphase target responses [6]. This is because some high-frequency energy of the impulse response is lost due to the shorter window length at high frequencies. Similarly, in quasi-logarithmic filter design methods a high-frequency attenuation can be observed for non-minimumphase targets. Therefore, the target response is often made minimum-phase before filter design, which is acceptable in several applications, where only magnitude modeling or equalization is needed.

However, some cases require precise phase modeling to match the time-domain structure of the target response. For example, in physics-based sound synthesis the radiation of the instrument body is often modeled by a filter, designed from measured impulse responses. Making the body response minimum-phase is not a good option because this eliminates the slow attack of the response and destroys its reverberant character, leading to an altered sound [7]. On the other hand, when using the original (non-minimumphase) target, high-frequency attenuation can be observed as mentioned above, which is again undesirable.

This is illustrated in Fig. 1 for piano soundbaord response modeling. The measured transfer function is displayed in Fig. 1 (a), upper pane. First, a 400th order parallel filter [1] [10] is designed from the original (complex) target response, displayed by Fig. 1 (b). While the filter follows the target nicely at low and mid frequencies, the high frequencies are attenuated due to the frequency-dependent windowing effect of the filter design. This can be overcome by a magnitude-only design, where the parallel filter

[1]Note that in this paper logarithmic pole positioning used for the parallel filter design examples, where the pole frequencies are spread uniformly on the logarithmic frequency scale. This technique is used here because its simplicity, more complex pole positioning methods are available in [8, 9] which generally result in more accurate modeling for the same filter order.

is designed using the minimum-phase version of the target, shown by Fig. 1 (b).

The same responses can be observed in the time-domain in Fig. 1 lower pane. The target impulse response is displayed by Fig. 1 (a), which is followed quite well for the first few hundred samples by the parallel filter designed using the original specification (b). It can also be seen that as time increases, the high-frequency content of the filter impulse response decreases due to the frequency-dependent windowing effect, which was already observed as a high-frequency attenuation in Fig. 1 (b) upper pane. While minimum-phase filter design corrects the high-frequency attenuation, the filter impulse response displayed in Fig. 1 (c) has nothing similar to the original. In particular, most of the energy is concentrated in the beginning of the response, as opposed to the slow attack of the measured impulse response.

In frequency-dependent windowing the high-frequency attenuation is overcome by correcting the smoothed magnitude to match that of the power-smoothed response [6]. Here the goal is to provide a similar effect in filter design. For that, two iterative algorithms are presented in the next section.

3. MAGNITUDE-PRIORITY FILTER DESIGN

3.1. Iterative magnitude update

In this technique, the magnitude of the target response is updated iteratively, while the phase specification is unchanged. The steps of the procedure are as follows:

1. Design a filter $H(\vartheta)^{(0)}$ for the original complex (non-minimumphase) target $H_{\mathrm{t}}(\vartheta)^{(0)} = H_{\mathrm{t},0}(\vartheta)$ by the complex filter design method of your choice. Set $i = 1$.

2. Compute a new specification by correcting the *magnitude* of the target used in step $(i-1)$ as

$$H_{\mathrm{t}}(\vartheta)^{(i)} = \frac{|H_{\mathrm{t},0}(\vartheta)|}{|H(\vartheta)^{(i-1)}|} H_{\mathrm{t}}(\vartheta)^{(i-1)} \quad (1)$$

3. Design a new filter $H(\vartheta)^{(i)}$ for the updated target $H_{\mathrm{t}}(\vartheta)^{(i)}$.

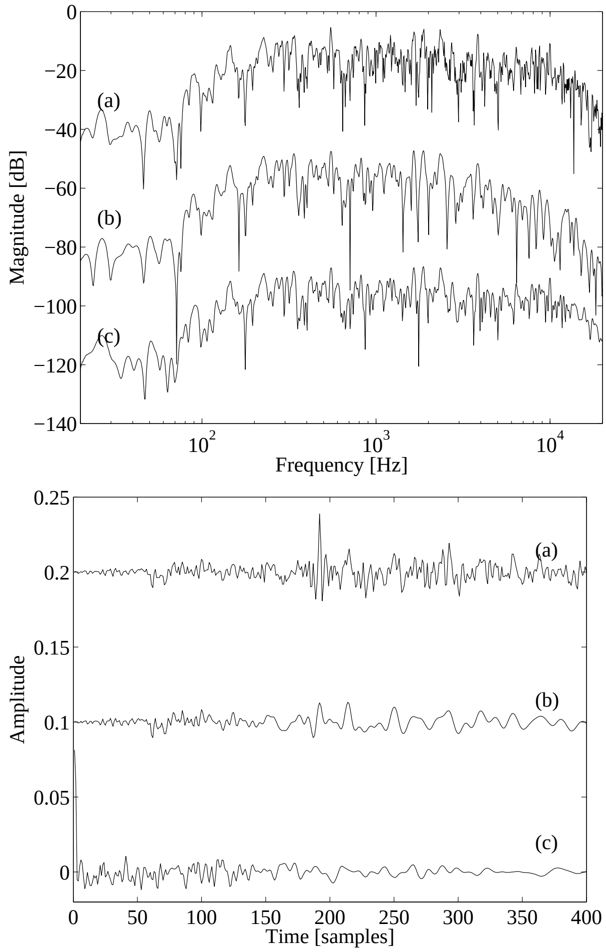

Fig. 1: Designing a 400th order parallel filter from a piano soundboard response: (a) target response, (b) the response of the filter designed from the original (complex) specification, and (c) the response of the filter designed from the minimum-phase version of the target. The upper pane shows magnitude responses, while the lower pane displays impulse responses. The curves are offset for clarity.

4. Increase i and go to step 2.

The iteration is repeated five-ten times.

The core of the algorithm is step 2, where the magnitude response of the target is multiplied by the ratio of the original target and previous filter magnitude responses, practically meaning that the magnitude of the target response is increased at those frequencies where the filter has an undesirable attenuation. In this way we are "pushing harder" at those frequencies where the magnitude matching is inadequate. While not indicated in the steps above for clarity, note that both the target $|H_{\mathrm{t},0}(\vartheta)|$ and filter $|H(\vartheta)^{(i-1)}|$ magnitude responses are fractional-octave smoothed before their ratio is computed by Eq. (1), that is, only an overall correction is done. Without smoothing, wherever the filter response $|H(\vartheta)^{(i-1)}|$ tends to zero, the ratio $|H_{\mathrm{t},0}(\vartheta)|/|H(\vartheta)^{(i-1)}|$ goes to infinity, heavily biasing the filter design. In the examples of Sec. 4 third-octave smoothing is used, but the resolution of smoothing is not critical.

3.2. Iterative phase update

A simpler solution that does not require the smoothing of transfer functions is based on the iterative update of the target phase. The steps are as follows:

1. Design a filter $H(\vartheta)^{(0)}$ for the original complex (non-minimumphase) target $H_{\mathrm{t}}(\vartheta)^{(0)} = H_{\mathrm{t},0}(\vartheta)$ by the complex filter design method of your choice. Set $i = 1$.

2. Compute a new specification by updating the *phase* of the target as

$$H_{\mathrm{t}}(\vartheta)^{(i)} = |H_{\mathrm{t},0}(\vartheta)| e^{j\varphi\{H(\vartheta)^{(i-1)}\}} \qquad (2)$$

3. Design a new filter $H(\vartheta)^{(i)}$ for the updated target $H_{\mathrm{t}}(\vartheta)^{(i)}$.

4. Increase i and go to step 2.

The iteration is repeated five-ten times.

The main idea of the method is in step 2, where the phase of the target response is updated to match the phase of the filter $\varphi\{H(\vartheta)^{(i-1)}\}$ designed in the previous step. This means that the filter design algorithm is given a phase response that it can easily follow, thus, it can focus on the magnitude response.

Note that the steps of the method are similar to the magnitude-only method of [11, 12], but here the trick is that the iteration is started from the *original* target, and not from an altered one (e.g., minimum-phase version).

One can see that the magnitude update and phase update methods have quite a different approach: the magnitude update method modifies the magnitude specification to force the filter to match the magnitude, while the phase update variant gives an "easy phase" so that the algorithm can focus on the magnitude. However, they result in very similar filters, as will be seen in the examples in Sec. 4. A common property of the two algorithms is that no target update is performed whenever the first filter matches both phase and magnitude.

4. DESIGN EXAMPLES

4.1. Loudspeaker modeling

The magnitude and phase responses of a small two-way loudspeaker are displayed in Fig. 2 by thin lines in the upper and lower panes, respectively. First, a 20th order parallel filter [3] with logarithmic pole positioning is designed from the complex loudspeaker response (a). It can be seen in the upper pane that the magnitude response of the filter (thick line) follows that of the target (thin line) quite well up to around 10 kHz, where the filter response gets attenuated due to the shorter effective window length of the filter design. On the other hand, the phase response (see lower pane of Fig. 2) is modeled quite well. When the parallel filter is designed from the minimum-phase version of the target, the high-frequency attenuation is avoided (see (b) in upper pane), however, the phase response differs from the original. Note that a linear term corresponding to a constant delay has been added to the phase response coming from the minimum-phase design to best match the target in the middle frequencies. Without this, the target and filter phase would differ even more, since minimum-phase design also removes the constant delay of the target response.

The results of the two variants of the new method are presented in Fig. 2 (c) and (d). It can be seen that

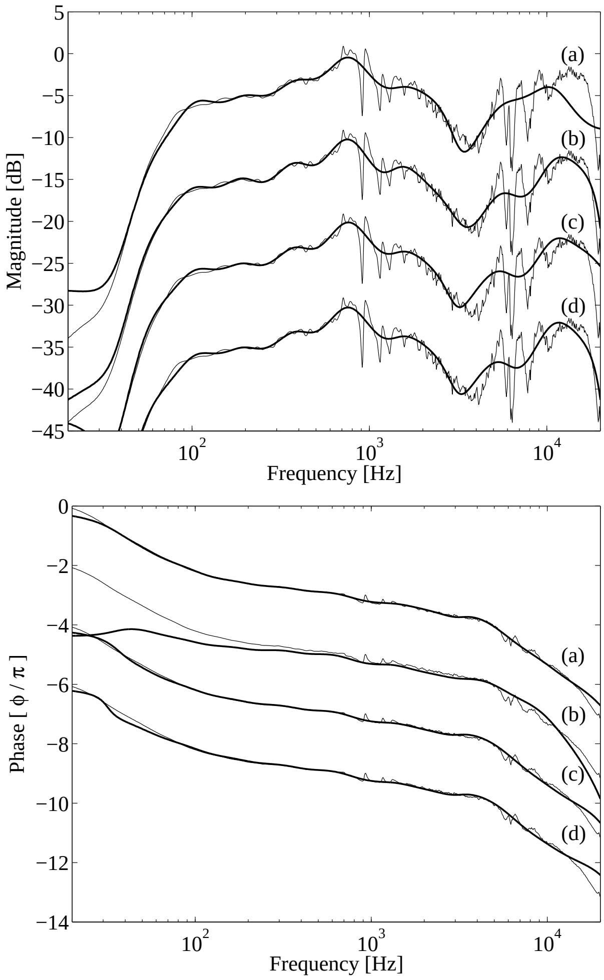

Fig. 2: Loudspeaker modeling using a 20th order parallel filter: (a) straightforward (complex) filter design using the original target, (b) magnitude-only filter design using the minimum-phase version of the target, (c) magnitude-priority filter design by the magnitude update method, and (d) magnitude-priority filter design by the phase update method. The thin lines show the target magnitude and phase responses, while the thick lines show the magnitude and phase responses of the designed filters. The curves are offset for clarity.

both the magnitude update (c) and phase update (d) methods result in a magnitude response that is similar to what is achieved by the minimum-phase design (b), while the phase response is only slightly less accurately modeled compared to the straightforward (complex) filter design (a). In this particular example, the magnitude update (c) results in a bit more accurate phase response, while the phase update (d) leads to a somewhat more accurate magnitude response. However, in general the two methods perform quite similarly, and provide a sensible alternative compared to the complex and magnitude-only filter designs.

4.2. Piano soundboard modeling

Here 400th order parallel filters are designed from the same piano soundboard response which was used in Fig. 1. It can be seen in the upper pane of Fig. 3 that both the magnitude update (b) and phase update (c) methods correct the high-frequency attenuation which was present in complex filter design (shown in Fig. 1 (b) upper pane). When looking at the time-domain responses in the lower pane of Fig. 3, the filter impulse responses (b) and (c) does not look very similar to the original (a). This is because the high-frequency attenuation can only be counteracted by increasing the high-frequency content of the impulse response in the early part of the filter response.

However, when the same responses are displayed after the high-frequency content above 5 kHz is removed, we see a different behavior in Fig. 4. In this case the magnitude update (c) and phase update (d) methods follow the target impulse response (thin lines) just as well as the straightforward filter design (a), but without the high-frequency attenuation that was visible in the frequency response in Fig. 1 (b) upper pane. The nice time-domain performance implies that the phase is well matched in this frequency range. On the other hand, it is no surprise that the low-pass filtered impulse response of the minimum-phase filter is displayed by (c) is very different from the lowpass-filtered target, since the target was non-minimumphase.

The similar performance of the complex and magnitude-priority methods below 5 kHz can be explained as follows. In this particular case the magnitude-priority methods model the time-domain

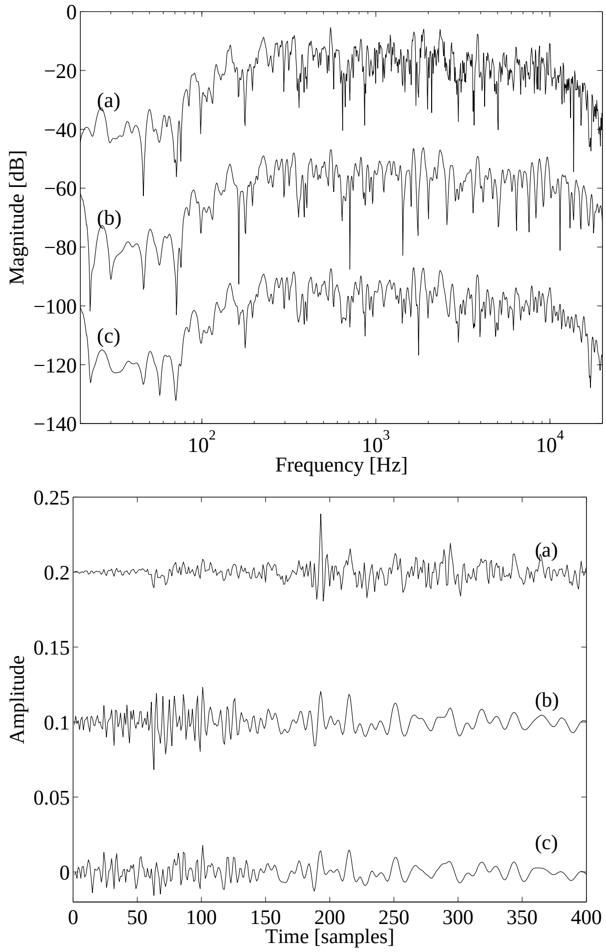

Fig. 3: Designing a 400th order parallel filter to a piano soundboard response: (a) target response, (b) the response of the filter designed by the magnitude update method, and (c) the response of the filter designed by the phase update method. The upper pane shows magnitude responses, while the lower pane displays imulse responses. The curves are offset for clarity.

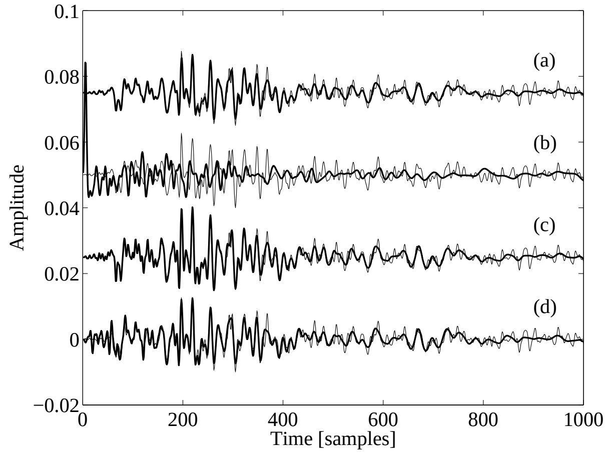

Fig. 4: Designing 400th order parallel filters to a piano soundboard response. The low-pass filtered (f_c = 5 kHz) impulse responses of the (a) parallel filter designed from the original (complex) specification, (b) the parallel filter designed from the minimum-phase version of the target, (c) the filter designed by the magnitude update method, and (d) the filter designed by the phase update method. Thick lines show the filter impulse responses, while the thin lines display the impulse response of the target. The curves are offset for clarity.

response quite precisely up to 5 kHz, since this is the frequency range where also the straightforward (complex) filter design was able to provide a good match. Hence, the filter specification is not modified below 5 kHz during the iterative updates. Above 5 kHz the straightforward filter design cannot provide precise modeling due to its shorter effective window length, and thus the magnitude-priority methods have to increase the high-frequency content in the early part of the filter response (visible in Fig. 3 (b) and (c) lower pane) to provide a good match in magnitude.

Figure 5 shows how the above filters perform in a real-life situation when used as a soundboard model as a part of a physics-based piano model. In the model the output of a physics-based piano string model [13] is fed to the parallel filters. The reference case is when the string signal is convolved by the measured piano soundboard response directly, displayed in Fig. 5 (a). One can see that both

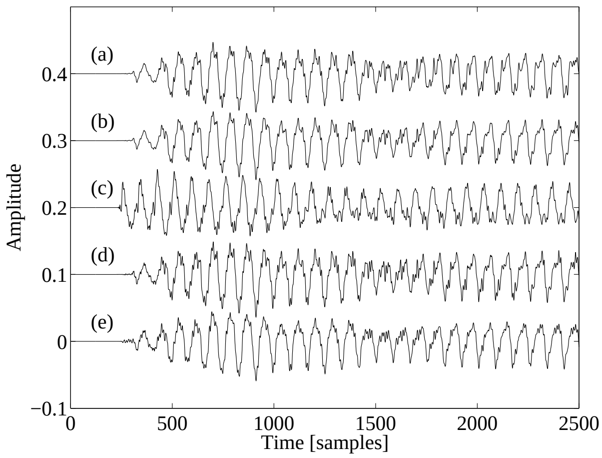

Fig. 5: A synthesized C_5 ($f_0 = 524$ Hz) piano tone computed by (a) convolving a synthesized string signal with the measured soundboard impulse response, (b) filtering by the parallel filter designed from the complex target, (c) filtering by the parallel filter designed from the minimum-phase target, (d) filtering by the parallel filter designed using the magnitude update method, and (e) filtering by the parallel filter designed by the phase update method. The curves are offset for clarity.

the complex method (b) and the magnitude-priority methods (d) and (e) follow the reference case (a) quite well. In particular, the relatively slow attack of the note is well captured. On the other hand, the magnitude-only method using a minimum-phase filter results in a sharper attack (c), which is undesirable.

By looking at Fig. 5 we might conclude that the straightforward complex method (a) performs just as well as the magnitude-priority methods, but we already know from Fig. 1 (b) upper pane that the complex method results in a high-frequency attenuation. While this attenuation is not visible in the time-domain signal of Fig. 5, it is in the spectral domain, and indeed results in a duller sound. As a result, the use of the magnitude-priority method is preferable, since it both models the spectral coloration of the soundboard adequately and preserves the characteristic attack of the piano sound.

5. CONCLUSION

Most filter design algorithms aim to match the complex transfer function (or, equivalently, the impulse response) of the target. Another common option is magnitude-only filter design. This latter is used either if only the magnitude response of the target is specified, or when the matching of the complex transfer function would not be possible due to the low filter order. This paper has presented a third type of design methodology, called magnitude-priority filter design. The new method is a useful alternative to magnitude-only design, since it matches the complex transfer function (thus, magnitude and phase) whenever possible, while in those frequency regions where this cannot be done, it follows the magnitude of the target just as well as the magnitude-only method.

This method is especially useful in audio applications because the quasi-logarithmic frequency-resolution filter design methods (e.g., warped, Kautz, and parallel filters) result in a frequency-dependent windowing of the impulse response leading to a high frequency attenuation, which must be corrected. Instead of sacrificing the phase response (or, equivalently, the time-domain structure of the impulse response) in the full frequency range by the use of a magnitude-only design, the magnitude-priority design will maintain precise magnitude and phase modeling up to the frequency it can, and it will “switch to magnitude-only mode” only where it is necessary.

The method has been demonstrated by using parallel filter designs, but since it does not make any assumption on the core filter design algorithm used in the iterations, it can be used in combinations with any technique which is based on the matching of the complex transfer function, like standard FIR, IIR, warped, or Kautz filters. While here only filter design examples were presented, the method can also be used for the equalization of transfer functions, e.g., for the equalization of loudspeaker responses.

6. ACKNOWLEDGEMENT

This work has been supported by the Bolyai Scholarship of the Hungarian Academy of Sciences.

7. REFERENCES

[1] Aki Härmä, Matti Karjalainen, Lauri Savioja, Vesa Välimäki, Unto K. Laine, and Jyri Huopaniemi. Frequency-warped signal processing for audio applications. *J. Audio Eng. Soc.*, 48(11):1011–1031, Nov. 2000.

[2] Tuomas Paatero and Matti Karjalainen. Kautz filters and generalized frequency resolution: Theory and audio applications. *J. Audio Eng. Soc.*, 51(1–2):27–44, Jan./Feb. 2003.

[3] Balázs Bank. Perceptually motivated audio equalization using fixed-pole parallel second-order filters. *IEEE Signal Process. Lett.*, 15:477–480, 2008. URL: `http://www.acoustics.hut.fi/go/spl08-parfilt`.

[4] Matti Karjalainen, Tuomoas Paatero, John N. Mourjopoulos, and Panagiotis D. Hatziantoniou. About room response equalization and dereverberation. In *Proc. IEEE Workshop Appl. of Signal Process. to Audio and Acoust.*, pages 183–186, New Paltz, NY, USA, Oct. 2005.

[5] Balázs Bank. Audio equalization with fixed-pole parallel filters: An efficient alternative to complex smoothing. In *Proc. 128th AES Conv., Preprint No. 7965*, London, UK, May 2010.

[6] Panagiotis D. Hatziantoniou and John N. Mourjopoulos. Generalized fractional-octave smoothing for audio and acoustic responses. *J. Audio Eng. Soc.*, 48(4):259–280, Apr. 2000.

[7] Matti Karjalainen, Vesa Välimäki, Heikki Räisänen, and Harri Saastamoinen. DSP equalization of electret film pickup for acoustic guitar. *J. Audio Eng. Soc.*, 48(12):1183–1193, 2000.

[8] Balázs Bank and Germán Ramos. Improved pole positioning for parallel filters based on spectral smoothing and multi-band warping. *IEEE Signal Process. Lett.*, Mar. 2011. In press.

[9] Balázs Bank. Warped IIR filter design with custom warping profiles and its application to room equalization. In *Proc. 130th AES Conv., Preprint No. 7965*, London, UK, May 2011.

[10] Balázs Bank. Direct design of parallel second-order filters for instrument body modeling. In *Proc. Int. Computer Music Conf.*, pages 458–465, Copenhagen, Denmark, Aug. 2007. URL: `http://www.acoustics.hut.fi/go/icmc07-parfilt`.

[11] Leland B. Jackson. Frequency-domain Steiglitz-McBride method for least-squares filter design, ARMA modeling, and periodogram smoothing. *IEEE Signal Process. Lett.*, 15:49–52, 2008.

[12] Balázs Bank. Logarithmic frequency scale parallel filter design with complex and magnitude-only specifications. *IEEE Signal Process. Lett.*, 18(2):138–141, Feb. 2011.

[13] Balázs Bank, Federico Avanzini, Gianpaolo Borin, Giovanni De Poli, Federico Fontana, and Davide Rocchesso. Physically informed signal-processing methods for piano sound synthesis: a research overview. *EURASIP J. on Applied Signal Process.*, 2003(10):941–952, Sep. 2003.