Audio Engineering Society
Convention Paper 7965
Presented at the 128th Convention
2010 May 22–25 London, UK

Audio Equalization with Fixed-Pole Parallel Filters: An Efficient Alternative to Complex Smoothing

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

Recently, the fixed-pole design of parallel second-order filters has been proposed to accomplish arbitrary frequency resolution similarly to Kautz filters, at $2/3$ of their computational cost. This paper relates the parallel filter to the complex smoothing of transfer functions. Complex smoothing is a well-established method for limiting the frequency resolution of audio transfer functions for analysis, modeling, and equalization purposes. It is shown that the parallel filter response is similar to the one obtained by complex smoothing the target response using a hanning window: a $1/\beta$ octave resolution is achieved by using $\beta/2$ pole pairs per octave in the parallel filter. Accordingly, the parallel filter can be either used as an efficient implementation of smoothed responses, or, it can be designed from the unsmoothed responses directly, eliminating the need of frequency-domain processing. In addition, the theoretical equivalence of parallel filters and Kautz filters is developed, and the formulas for converting between the parameters of the two structures are given. Examples of loudspeaker-room equalization are provided.

1. INTRODUCTION

Audio equalization using DSPs has been a subject of research for more than two decades. It generally means the correction of the magnitude (and sometimes the phase) response of an audio chain. Typical examples include loudspeaker equalization based on anechoic measurements [1, 2, 3], or the correction of loudspeaker-room responses [4, 5, 6, 7, 8]. Because the systems to be equalized are typically of very higher order (e.g., due to the high modal density in room responses), the direct inversion of the transfer function is usually not practical. As the final judge in sound quality is the human ear, it is more efficient to equalize only those aspects that lead to an audible error. A typical approach is to take into ac-

count the quasi-logarithmic frequency resolution of the human auditory system during equalizer design.

Besides efficiency and perceptual aspects, there is also a physical reason for applying logarithmic or logarithmic-like frequency resolution. Namely, an audio system often has multiple outputs, like multiple listening positions in a room, and the equalizer should maintain or improve the sound quality at all positions. Transfer functions measured at different positions in space have more similarity at low frequencies than at high frequencies, due to the different wavelengths of sound. An overly precise correction at high frequencies for one measurement position usually worsens the response at other points in space. Accordingly, the direct inversion of measured transfer functions not only wastes computational resources, but also leads to unacceptable results [5, 6].

This paper demonstrates that fixed-pole parallel filters can be efficiently used for the modeling or equalization of audio systems, as they posses the beneficial properties of complex smoothing, while require relatively little processing power both for filtering and for parameter estimation. The paper first reviews fractional-octave smoothing in Sec. 2, then covers the related warped and Kautz filters in Sec. 3. Section 4 outlines the theory of parallel filters and proves their equivalence with Kautz filters. Section 5 relates parallel filter design to complex smoothing, and Sec. 6 presents loudspeaker-room equalization examples and comparison. Finally, Sec. 7 gives practical implications and Sec. 8 concludes the paper.

2. FRACTIONAL-OCTAVE SMOOTHING OF TRANSFER FUNCTIONS

The quasi-logarithmic frequency resolution of human hearing is also reflected in how transfer functions are displayed in the audio field. From the earliest times, a logarithmic frequency scale is used, and often the magnitude response is smoothed at a fractional-octave (e.g., third octave) resolution. The motivation behind fractional-octave smoothing is that the original transfer function is too detailed for visual evaluation and the smoothed version gives a good estimate of the perceived timbre. While this practice stems from analog signal analyzers, typically all current digital spectrum analyzers for audio offer this option.

While traditionally, smoothing has only been applied for the magnitude response of audio systems, the complex alternative has also appeared, allowing the reconstruction of a smoothed impulse response via IFFT [9]. Smoothing of the complex transfer function is in practice done by convolving it with a smoothing function, thus, it is equivalent to multiplying the impulse response with a time-domain window function.

Note that in complex smoothing, the smoothing function is chosen to be a real (zero phase) function [9], which has an important consequence that is not very much discussed in the literature. Since the frequency-domain smoothing function is real, the corresponding time window will be symmetric around $t = 0$, that is, $w(-t) = w(t)$. In continuous time, this should not pose any problems. However, in practice smoothing is done on a discrete-time sequence $h(n)$ of finite length N, and in this case the $t < 0$ part of the time domain window appears at the end of the window sequence $w(N - n) = w(n)$, (see Fig. 2 (b) of [9]). As a result, the $n > N/2$ part of $h(n)$ is treated as negative times by the discrete-time smoothing operation, [1] thus, it should be zero for causal impulse responses. Therefore, the initial half window $w_0(n)$ applied to $h(n)$ before smoothing should be of length $N/2$, and not N as suggested in [9].

When the width of the frequency-domain smoothing function depends on frequency, the operation equals to multiplying the impulse response with window whose length is frequency dependent. In practice, this means that the original impulse response is windowed to shorter length at high frequencies compared to the low ones, which also has some connections to how reflections are processed in our auditory system [9]. Frequency-dependent signal windowing, which is equivalent to transfer function smoothing as noted above, has also been proposed in [10]. Naturally, not only fixed fractional-octave (logarithmic), but arbitrary smoothing resolution can be applied, including those corresponding to Bark or ERB scales [9].

Note that the magnitude of the complex-smoothed transfer function will differ from that of the traditional magnitude or power smoothing, since some energy will be "lost" in the high frequencies due to the shorter corresponding time window [9]. This effect can be eliminated if needed by the use of "equivalent complex smoothing" [9], where the magnitude of the smoothed transfer function is corrected to match that of the power-smoothed

[1] This is similar how DFT represents negative frequencies in the upper half of the data set.

version. Alternatively, the magnitude and phase of the transfer function can be smoothed separately [11].

Besides signal analysis, complex-smoothed transfer functions have been successfully applied for loudspeaker-room response equalization, where the complex-smoothed impulse response is used for FIR inverse filter design [12]. In addition, some kind of transfer function smoothing (mostly magnitude smoothing) is applied in most room equalization systems as a preprocessing step before filter design (see, e.g., [3, 5, 8]), to avoid the problems of direct inversion mentioned in the Introduction.

3. WARPED AND KAUTZ FILTERS

Traditional FIR and IIR filter design techniques or system identification methods provide a linear frequency resolution, as opposed to the quasi-logarithmic resolution of hearing. Therefore, in audio, often specialized filter design methodologies are used. While there are many different techniques that take into account the frequency resolution of hearing, only those are addressed here that have a direct connection to complex smoothing, or, equivalently, to frequency dependent windowing.

3.1. Warped filters

The most commonly used perceptually motivated design technique is based on frequency warping [13, 14]. The basic idea of warped filters is that the unit delay z^{-1} in the traditional FIR or IIR filters is replaced by an allpass filter

$$z^{-1} \leftarrow D(z) = \frac{z^{-1} - \lambda}{1 - \lambda z^{-1}}. \quad (1)$$

By a particular choice of the λ parameter, it is possible to match the Bark or ERB scale closely [15].

The design of warped filters starts with warping the target impulse response $h_{\mathrm{t}}(n)$, e.g., by the use of an allpass chain, which can be considered as the frequency-dependent resampling of the impulse response. Then, warped FIR (WFIR) filters can be obtained by truncating or windowing the warped target response $\tilde{h}_{\mathrm{t}}(n)$. It follows directly from this design principle that WFIR filters perform frequency-dependent windowing [10], because the traditional, fixed-length windowing operation is embedded in a frequency-dependent resampling and its inverse operation. Accordingly, WFIR filters have been used for magnitude equalization of room responses with similar results compared to FIR filters obtained by complex smoothing, at a significantly lower computational cost [6]. However, the choice of the frequency profile of the "built in smoothing function" is limited in the case of WFIR filters, as the frequency transformation is controlled by only one parameter λ.

This is illustrated in Fig. 1, where the minimum-phase version of a loudspeaker-room response (a) is modeled by 32nd order warped FIR filters with $\lambda = 0.55; 0.85; 0.95$ parameters, displayed in (b), (c), and (d), respectively. While (a) solid line displays the target response after third-octave smoothing, [2] the WFIR filters are designed using the unsmoothed target response, to demonstrate the smoothing behaviour of WFIR filters. It can be seen that the allocation of the frequency resolution is controlled by the λ parameter, and lower values lead to better high frequency resolution, while larger ones increase the resolution at low frequencies. It can be also seen that none of the λ parameters distribute the resolution evenly on the logarithmic frequency scale.

Note that for warped IIR (WIIR) filters no direct connections with complex smoothing can be made.

3.2. Kautz filters

Kautz filters can be seen as the generalization of WFIR filters, where the allpass filters in the chain are not identical [17, 18]. As a result, the frequency resolution can be allocated arbitrarily by the choice of the filter poles. The Kautz structure is a linear-in-parameter model, where the basis functions are the orthonormalized versions of decaying exponentials. The transfer function of the Kautz filter is:

$$H(z) = \sum_{k=0}^{N} w_k G_k(z) = \sum_{k=0}^{N} w_k \left(\frac{\sqrt{1 - p_k p_k^*}}{1 - p_k z^{-1}} \prod_{j=0}^{k-1} \frac{z^{-1} - p_j^*}{1 - p_j z^{-1}} \right), \quad (2)$$

[2] Note that there is no generally accepted standard how third-octave smoothing should be performed. In this paper, third-octave smoothing corresponds to convolving the transfer function with a hanning window having the full width of 2/3-octave (i.e., its half width, or, the distance of its 0.5 points is third octave). This complies with the results of analog third-octave analyzers, where the half-power points of the bandpass filters have a third-octave distance [16].

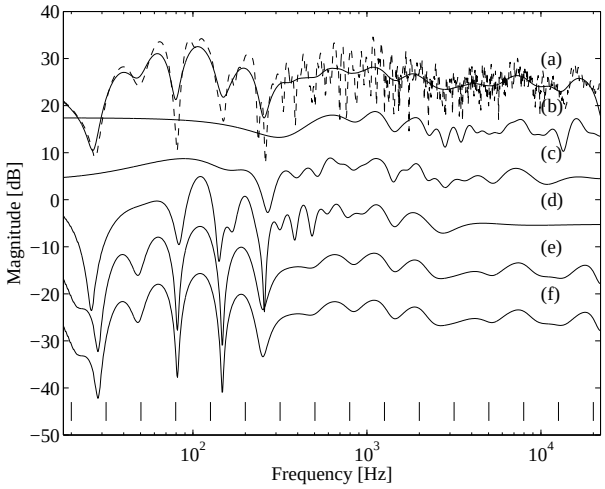

Fig. 1: Loudspeaker-room response modeling comparison: (a) third-octave smoothed target response (dashed line without smoothing), (b)-(d) 32nd order WFIR filters with $\lambda = 0.55; 0.85; 0.95$ parameters, (e) Kautz filter with 16 logarithmically spaced pole pairs (filter order is 32), and (f) parallel filter with the same pole set as for the Kautz filter. The vertical lines indicate the pole frequencies of the Kautz and parallel filters.

where $G_k(z)$ are the orthonormal Kautz functions determined by the pole set p_k, and p_k^* are the complex conjugate of p_k. The advantage of the orthonormality of $G_k(z)$ functions is that the weights w_k can be determined from the target response $h_t(n)$ by a scalar product

$$w_k = \sum_{n=1}^{\infty} g_k(n) h_{\mathrm{t}}(n), \tag{3}$$

where $g_k(n)$ are the inverse z transform of $G_k(z)$.

It is impractical to implement Kautz filters by a series of complex first-order allpass filters as in Eq. (2), and combining the complex pole pairs to second-order sections yields lower computational complexity [17]. However, the combined cascade-parallel nature of the filter still requires more computation compared to filters implemented in direct or cascade form with the same filter order.

For determining the poles p_k of the Kautz filter, several methods are discussed in [17], including some iterative techniques. For our purposes, the most interesting is the one that sets the poles according to the required resolution, e.g., by applying a logarithmic pole distribution. This is illustrated in in Fig. 1 (e), displaying the frequency response of a Kautz filter designed to match the unsmoothed target response of Fig. 1 (a), dashed line. It can be seen that by placing 2/3 poles per octave (vertical lines in Fig. 1), the resulting filter response approximates the third-octave smoothed transfer function, Fig. 1 (a) solid line. It is important to stress that the Kautz filter was designed from the unsmoothed response, thus, it performs smoothing "automatically". The present paper provides a theoretical explanation to this phenomenon, after proving the equivalence of Kautz and parallel filters in Sec. 4.4.

4. THE PARALLEL FILTER

Recently, a fixed-pole design method has been introduced for parallel second-order filters [19, 20]. It has been shown that effectively the same results can be achieved by the parallel filters as with Kautz filters for the same filter order, at 2/3 of their computational cost.

Implementing IIR filters in the form of parallel second-order sections has been used traditionally because it has better quantization noise performance and the possibility of code parallelization. The parameters of the second-order sections are usually determined from direct form IIR filters by partial fraction expansion [21]. In contrast, here the poles are set to a predetermined (e.g., logarithmic) frequency scale, leaving the zeros as free parameters for optimization. As we shall see later, the motivation for fixing the poles is to gain control over the frequency resolution of the design.

4.1. Problem formulation

Every transfer function of the form $H(z^{-1}) = B(z^{-1})/A(z^{-1})$ can be rewritten in the form of partial fractions:

$$H(z^{-1}) = \sum_{k=1}^{K} c_k \frac{1}{1 - p_k z^{-1}} + \sum_{m=0}^{M} b_m z^{-n}, \tag{4}$$

where p_k are the poles, forming either conjugate pairs or real valued, if the system has a real impulse response. The second sum in Eq. (4) is the FIR part of order M. If the order of $A(z^{-1})$ and $B(z^{-1})$ is the same, then it reduces to a constant coefficient b_0. Note that Eq. (4) is valid only if no multiple poles are present. In the case of pole multiplicity, terms of higher order also appear.

Now let us assume that we are trying to fit the filter $H(z^{-1})$ to a target response, but the poles of our IIR

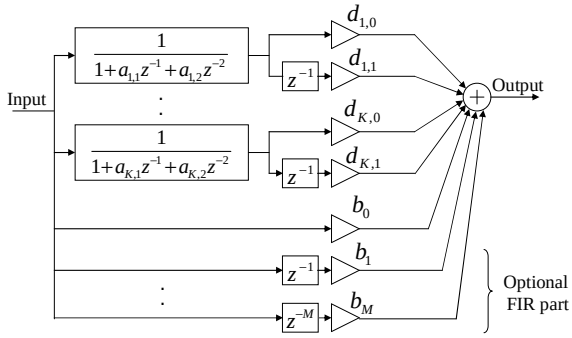

Fig. 2: Structure of the parallel second-order filter.

filter are also known. In this case Eq. (4) becomes linear in parameters c_k and b_m, thus, they can be estimated by a simple least squares algorithm to match the required response.

The resulting filter can be implemented directly as in (4), forming parallel first-order complex filters, and the estimation of the parameters can be carried out as described in [19]. However, it is more practical to combine the complex pole pairs to a common denominator. This results in second-order sections with real valued coefficients, which can be implemented more efficiently. Those fractions of (4) that have real poles can be combined with other real poles to form second-order IIR filters, yielding a canonical structure. Thus, the transfer function becomes

$$H(z^{-1}) = \sum_{k=1}^{K} \frac{d_{k,0} + d_{k,1} z^{-1}}{1 + a_{k,1} z^{-1} + a_{k,2} z^{-2}} + \sum_{m=0}^{M} b_m z^{-n} \tag{5}$$

where K is the number of second order sections. The filter structure is depicted in Fig. 2.

In the context of approximating complex smoothing, the pole frequencies f_k should be set according to the required frequency resolution. Accordingly, the poles of the parallel filter, p_k, are computed by the following formulas:

$$\vartheta_k = \frac{2\pi f_k}{f_s} \tag{6a}$$

$$p_k = e^{-\frac{\Delta\vartheta_k}{2}} e^{\pm j\vartheta_k}, \tag{6b}$$

where ϑ_k are the pole frequencies in radians given by the predetermined (e.g., logarithmic) analog frequency series f_k and the sampling frequency f_s. The bandwidth of the kth second-order section $\Delta\vartheta_k$ is computed from the neighboring pole frequencies

$$\begin{aligned} \Delta\vartheta_k &= \frac{\vartheta_{k+1} - \vartheta_{k-1}}{2} \quad \text{for } k = [2, .., K-1] \\ \Delta\vartheta_1 &= \vartheta_2 - \vartheta_1 \\ \Delta\vartheta_K &= \vartheta_K - \vartheta_{K-1} \end{aligned} \tag{7}$$

Equation (6b) sets the pole radii $|p_k|$ in such a way that the transfer functions of the parallel sections cross approximately at their -3dB point.

4.2. Filter design

First, we investigate how the parameters of the parallel filter can be estimated to match a desired filter response. The simplest way is to find the coefficients in the time domain. The impulse response of the parallel filter is given by

$$h(n) = \sum_{k=1}^{K} d_{k,0} u_k(n) + d_{k,1} u_k(n-1) + \sum_{m=0}^{M} b_m \delta(n-m) \tag{8}$$

where $u_k(n)$ is the impulse response of the transfer function $1/(1 + a_{k,1} z^{-1} + a_{k,2} z^{-2})$, which is an exponentially decaying sinusoidal function, and $\delta(n)$ is the discrete unit impulse.

Conceptually, filter design simply consists of creating a weighted sum of the exponentially decaying sinusoidal functions (and unit impulses if there is an FIR part) in such a way that the resulting signal best approximates the target impulse response $h_\mathrm{t}(n)$.

Because (8) is linear in parameters, it can be written in a matrix form:

$$\mathbf{h} = \mathbf{M}\mathbf{p} \tag{9}$$

where $\mathbf{p} = [d_{1,0}, d_{1,1}, \ldots d_{K,0}, d_{K,1}, b_0 \ldots b_M]^T$ is a column vector composed of the free parameters. The columns of the modeling signal matrix $\mathbf{M}$ contain the modeling signals, which are $u_k(n)$ and their delayed counterparts $u_k(n-1)$, and for the FIR part, the unit impulse $\delta(n)$ and its delayed versions up to $\delta(n-M)$. Finally, $\mathbf{h} = [h(0) \ldots h(N)]^T$ is a column vector composed of the resulting impulse response. The problem reduces to finding the optimal parameters $\mathbf{p}_\mathrm{opt}$ such that $\mathbf{h} = \mathbf{M}\mathbf{p}_\mathrm{opt}$ is closest to the target response $\mathbf{h}_\mathrm{t}$. If the error function is evaluated in the mean squares sense, the

optimum is found by the well known LS solution

$$\mathbf{p}_{\mathrm{opt}} = \mathbf{M}^{+}\mathbf{h}_{\mathrm{t}} \quad (10a)$$
$$\mathbf{M}^{+} = (\mathbf{M}^{H}\mathbf{M})^{-1}\mathbf{M}^{H} \quad (10b)$$

where $\mathbf{M}^{+}$ is the Moore-Penrose pseudoinverse, and $\mathbf{M}^{H}$ is the conjugate transpose of $\mathbf{M}$. Note that if the frequency resolution – thus the pole set and modeling matrix $\mathbf{M}$ – is fixed, the pseudo-inverse $\mathbf{M}^{+}$ can be precomputed and stored, so the parameter estimation reduces to a matrix multiplication according to Eq. (10a).

In Fig. 1 (f), a parallel filter design example is presented for the same loudspeaker-room target response as for warped and Kautz filters in Sec. 3. The same logarithmic pole set is used as for the Kautz filter. It can be seen that the parallel filter (f) results in the same filter response as the Kautz filter (e).

4.3. Direct equalizer design

Equalizing a system (such as a loudspeaker) by the parallel filter can be done by first inverting the system response and designing the parallel filter as outlined in the previous section. However, it is more practical to design the equalizer directly without inverting the system response [20], since this simplifies the design and avoids many problems presented by the direct inversion of the measured transfer function.

Designing an equalizer requires that the resulting response $h(n)$, which is the convolution of the equalizer response $h_{\mathrm{eq}}(n)$ and the system response $h_{\mathrm{s}}(n)$, is close to the target response $h_{\mathrm{t}}(n)$ (which can be a unit impulse, for example). By looking at Fig. 2, the basic idea of equalizer design is the following: imagine the system response $h_{\mathrm{s}}(n)$ is fed to the input of the parallel filter, and the weights $d_{k,0}$, $d_{k,1}$, and b_m should be set in such a way that the summed output best approximates the target impulse response $h_{\mathrm{t}}(n)$.

Accordingly, the output of the parallel filter is computed as

$$\begin{aligned} h(n) &= h_{\mathrm{eq}}(n) * h_{\mathrm{s}}(n) = \\ &\sum_{k=1}^{K} d_{k,0}u_k(n) * h_{\mathrm{s}}(n) + d_{k,1}u_k(n-1) * h_{\mathrm{s}}(n) + \\ &\sum_{m=0}^{M} b_m\delta(n-m) * h_{\mathrm{s}}(n) = \\ &\sum_{k=1}^{K} d_{k,0}s_k(n) + d_{k,1}s_k(n-1) + \sum_{m=0}^{M} b_m h_{\mathrm{s}}(n-m) \end{aligned} \quad (11)$$

where $*$ denotes convolution. The signal $s_k(n) = u_k(n) * h_{\mathrm{s}}(n)$ is the system response $h_{\mathrm{s}}(n)$ filtered by $1/(1 + a_{k,1}z^{-1} + a_{k,2}z^{-2})$. It can be seen that (11) has the same structure as (8). Therefore, the parameters $d_{k,0}$, $d_{k,1}$, and b_m can be estimated in the same way as presented in the previous section. Similarly, writing this in a matrix form yields

$$\mathbf{h} = \mathbf{M}_{\mathrm{eq}}\mathbf{p} \quad (12)$$

where the columns of the new signal modeling matrix $\mathbf{M}_{\mathrm{eq}}$ contain $s_k(n)$, $s_k(n-1)$, and the system response $h_{\mathrm{s}}(n)$ and its delayed versions up to $h_{\mathrm{s}}(n-M)$. Then, the optimal set of parameters is again obtained by

$$\mathbf{p}_{\mathrm{opt}} = \mathbf{M}_{\mathrm{eq}}^{+}\mathbf{h}_{\mathrm{t}} \quad (13a)$$
$$\mathbf{M}_{\mathrm{eq}}^{+} = (\mathbf{M}_{\mathrm{eq}}^{H}\mathbf{M})^{-1}\mathbf{M}_{\mathrm{eq}}^{H}. \quad (13b)$$

4.4. Equivalence with Kautz filters

While it was clear from simulations and theoretical reasoning [20] that Kautz and parallel filters result in the same filter response for the same pole set (see also Fig. 1), a formal proof has not been provided previously. The proof presented here is based on the partial fraction expansion of the Kautz basis functions. According to Eq. (2), the kth basis function of the Kautz filter is

$$G_k(z) = \frac{\sqrt{1-p_k p_k^*}}{1-p_k z^{-1}} \prod_{j=0}^{k-1} \frac{z^{-1}-p_j^*}{1-p_j z^{-1}}, \quad (14)$$

which is a kth order filter. In the case of no multiple poles, which is easily satisfied when the pole set is predetermined, Eq. (14) can be written in a partial fraction form

$$G_k(z^{-1}) = \sum_{i=1}^{k} c_{k,i} \frac{1}{1-p_i z^{-1}}, \quad (15)$$

where the k coefficients $c_{k,i}$ are found by the usual procedure of partial fraction expansion [21] in a closed form:

$$c_{k,i} = \sqrt{1 - p_k p_k^*} \prod_{j=1, j\neq i}^{k} \frac{1}{p_i - p_j} \prod_{j=1}^{k-1} (1 - p_j^* p_i). \quad (16)$$

By noting that the partial fraction form of Eq. (15) is the same as the complex form of the parallel filter Eq. (4) without the FIR part (M=0), it is clear that the Kautz basis functions can be reconstructed by the parallel filter exactly. As the Kautz filter response is the linear combination of the Kautz basis functions $G_k(z)$, it is straightforward to convert a Kautz filter into a parallel filter. If the parameters of the Kautz filter are ginven in a vector $\mathbf{w} = [w_1, \ldots, w_K]^T$, the parameter vector of the parallel filter $\mathbf{c} = [c_1, \ldots, c_K]^T$ can be obtained by the matrix multiplication

$$\mathbf{c} = \mathbf{K}\mathbf{w} \quad (17)$$

where the conversion matrix $\mathbf{K}$ is given as

$$K_{i,k} = \sqrt{1 - p_k p_k^*} \prod_{j=1, j\neq i}^{k} \frac{1}{p_i - p_j} \prod_{j=1}^{k-1} (1 - p_j^* p_i) \quad \text{for } i \leq k,$$
$$K_{k,i} = 0 \text{ for } i > k \quad (18)$$

Since $\mathbf{K}$ is tridiagonal, and for simple poles it does not have zero values in its diagonal, it is nonsingular. As a result, the inverse matrix $\mathbf{K}^{-1}$ can be computed that can be used to convert the parallel filter parameters to Kautz parameters ($\mathbf{w} = \mathbf{K}^{-1}\mathbf{c}$).

Basically, we have shown that the basis functions of the parallel and Kautz filters span the same approximation space, and converting between the two filters is merely a base change. Therefore, approximating a target response using any error norm (e.g., the L_2 norm in least-squares design) will lead to exactly the same filter response in both cases.

Besides its theoretical importance, the possibility of converting the Kautz parameters to the parallel filter parameters allows a computationally more efficient design of the parallel filter. Namely, first a Kautz filter is designed by the scalar product of Eq. (3), then the parameters are converted by Eqs. (17) and (18). While this seems to be conceptually more complicated, the number of required arithmetic operations is reduced compared to the LS design of Eq. (10), so it is a useful alternative for high (> 100) filter orders. Note that in the case of direct equalizer design of Sec. 4.3 this procedure does not provide any computational benefits compared to the LS design of Eq. (13), since for that case the scalar product of Eq. (3) cannot be used and also the Kautz filter has to be designed by a LS equation [18].

5. AN EFFICIENT ALTERNATIVE TO COMPLEX SMOOTHING

By observing the results of parallel and Kautz filters (see Fig. 1 (e) and (f)), it is apparent that the effect of filter design is similar to that of the fractional-octave complex-smoothing of transfer functions. However, the theoretical reasons for this similarity have remained unexplored. In this paper, the case of the parallel filter is investigated, but since it results in exactly the same filter response as the Kautz filter (as it was proven in Sec. 4.4), the observations can be generalized for the Kautz filter as well.

5.1. Linear pole distribution

We start our analysis with the simplest case, where the K poles of the parallel filter are distributed evenly on a circle of radius $R < 1$. After some algebra, it is relatively straightforward to demonstrate that the transfer function of the parallel filter with such a pole distribution is actually equivalent to a comb filter and a $(K-1)$th order FIR filter $B(z^{-1})$ in series:

$$\sum_{k=1}^{K} \frac{c_k}{1 - Re^{j2\pi k/K} z^{-1}} = B(z^{-1}) \frac{1}{1 - z^{-K} R^K} \approx B(z^{-1}). \quad (19)$$

In practice, $R^K \ll 1$, therefore, the transfer function can be approximated by the FIR part only. Since the coefficients of the FIR filter $B(z^{-1})$ can be computed as the linear combination of the parameters of the parallel filter c_k, designing the filter in the FIR form is equivalent to designing it in its original parallel form. When designing the parallel filter according to Eq. (10), the mean-squared error between the target impulse response $h_\mathrm{t}(n)$ and the filter response $h(n)$ is minimized. In the equivalent FIR design, the mean-squared error is minimal if the K coefficients of the FIR filter are chosen to be equal to the first K samples of the target response $h_\mathrm{t}(n)$. Therefore, the resulting impulse response $h(n)$ is the truncated version

of the target response $h_\mathrm{t}(n)$, which is equivalent to multiplying the target response with a rectangular window $w(n)$ of length K.

Note that this window is defined only for positive times $n \geq 0$ (it is a half window), in contrast to the symmetric window used in complex smoothing (see Sec. 2). Since for causal impulse responses $h(n) = 0$ for $n < 0$, we may think of extending the window $w(n)$ to negative times by setting $w(-n) = w(n)$, without influencing the result $h(n) = h_\mathrm{t}(n)w(n)$. This has the advantage that now the results will be directly comparable with those of complex smoothing.

Accordingly, designing a parallel filter with a linear pole distribution is equivalent to multiplying the target response with a symmetric rectangular window of total length $2K-1$. In the frequency domain, this corresponds to convolving the transfer function with a *sinc*-like (periodic sinc) function:

$$H(\vartheta) = H_\mathrm{t}(\vartheta) * \frac{\sin\left(\frac{2K-1}{2}\vartheta\right)}{\sin\left(\frac{1}{2}\vartheta\right)} \qquad (20)$$

which is clearly a form of transfer function smoothing. Actually, it corresponds to "filtering" the transfer function with an ideal lowpass filter. It can also be seen that the smoothing function does not depend on frequency, as expected from linear frequency resolution. Note that the main lobe of the *sinc*-like function in Eq. 20 approximates a hanning window quite closely. The width of the main lobe is $4\pi/(2K-1) \approx 2\pi/K$, therefore, the effect will be similar to smoothing the transfer function with a $2\pi/K$ wide symmetric hanning window, where $2\pi/K$ equals to the distance of pole frequencies $\Delta\vartheta$.

5.2. Stepwise linear pole distribution

Next, let us consider a more interesting case, when the pole density is different in the various regions of the frequency range, but it is is constant within each region. In this case, it is not possible to derive the smoothing function in a closed form, as in Sec. 5.1. Therefore, a different approach is taken.

Since the parameters of the parallel filter are determined by a linear LS design, the superposition principle holds. This means that if we decompose the target response $h_\mathrm{t}(n)$ as a sum of some test functions, the filter response $h(n)$ will be equal to the sum of filter responses designed for the test functions separately. Since we would like to gain some insight to the frequency-dependent nature of smoothing, a natural choice for such a test function is the basis function of the Fourier transform $e^{-j\vartheta_0 n}$, where ϑ_0 is the angular frequency of the complex exponential. In the frequency domain, this is equivalent to $\delta(\vartheta - \vartheta_0)$, which is a Dirac delta function at position ϑ_0. Accordingly, in the frequency domain, we are computing the "impulse response" of the smoothing operation, that is, we obtain the smoothing function directly.

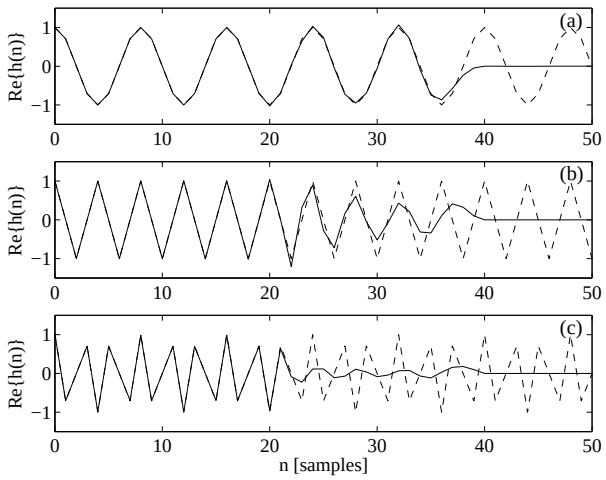

Fig. 3: Modeling a complex exponential $e^{-j\vartheta_0 n}$ with a parallel filter having stepwise linear frequency resolution: (a) $\vartheta_0 = 1/4\pi$, (b) $\vartheta_0 = 1/2\pi$, and (c) $\vartheta_0 = 3/4\pi$. The dashed line is the target response $e^{-j\vartheta_0 n}$ and the solid line is the resulting parallel filter impulse response. Only real parts of the signals are shown.

If the ovarlap of the basis functions of the parallel filter is not too large, our test function $e^{-j\vartheta_0 n}$ will be approximated by parallel sections whose center frequency is near to ϑ_0, while the contribution of the other sections will be negligible. Therefore, we expect that the width of the smoothing function in the frequency domain, and the length of the corresponding window function in the time domain will only depend on the *local* pole density near ϑ_0.

From Sec. 5.1 we may stipulate that if the the distance of the poles is $\Delta\vartheta$ in some frequency region, the "width" of the smoothing function in that region should be $\Delta\vartheta$, and signals in that frequency range should be windowed to a length $2\pi/\Delta\vartheta$ (or, the equivalent symmetric time window should have the length $4\pi/\Delta\vartheta - 1$).

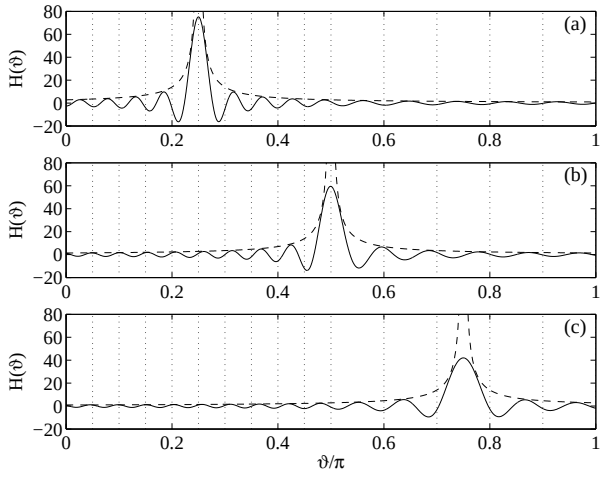

Fig. 4: Smoothing functions corresponding to the cases of Fig. 3 (a)-(c). The dotted vertical lines display the pole frequencies of the parallel filter. The dashed lines show the $1/|\vartheta - \vartheta_0|$ envelope of the transfer functions.

This is illustrated in Figs. 3 and 4, displaying a parallel filter design with 30 poles (15 pole pairs) around the unit circle. The pole frequencies are chosen in such a way that 20 poles are distributed evenly in the lower half of the frequency range $|\vartheta| \leq \pi/2$, while 10 poles are spread in the upper range $\pi/2 < |\vartheta| \leq \pi$. That is, the pole frequency distance $\Delta\vartheta_1 = 1/20\pi$ is the half in the low range compared to the high one $\Delta\vartheta_2 = 1/10\pi$ (see the dotted vertical lines in Fig. 4).

In Fig. 3 the target impulse responses $e^{-j\vartheta_0 n}$ are displayed by dashed lines, and the resulting parallel filter responses by solid lines. Note that the target and filter responses are complex, here only the real parts of the signals are shown, since the imaginary parts would show a similar behavior. Figure 3 (a) shows a case when the frequency of the exponential test function is in the high pole density region, while in (c) the frequency is in the low pole density region of the filter. As expected, the resulting impulse response (solid line) is "windowed" to a longer length in the first case compared to the second case. The theoretically computed half window length $2\pi/\Delta\vartheta$ is 40 and 20 for (a) and (c), which is in a good agreement with what can be observed in practice. Figure 3 (b) displays an intermediate case when the frequency of the test signal is exactly at the boundary of the two pole density regions. This results in a more mild windowing, where the window length is somewhere in between the (a) and (c) cases.

The same phenomenon can also be observed in the frequency domain in Fig. 4, for the same cases. The solid lines display the transfer functions of the resulting filters trying to approximate the test function $e^{-j\vartheta_0 n}$, again with (a) $\vartheta_0 = 1/4\pi$, (b) $\vartheta_0 = 1/2\pi$, and (c) $\vartheta_0 = 3/4\pi$. Note that the frequency responses were computed by extending the parallel filter responses to negative times $h(-n) = h^*(n)$, to comply with the symmetric windows used in complex smoothing. Accordingly, the frequency responses displayed in Fig. 4 are real (zero phase) functions, and can be directly compared by the smoothing windows used in complex smoothing. In Fig. 4, the dotted vertical lines show the pole frequencies of the parallel filter. It can be seen in (a) and (c), that the width of the main lobe equals to the pole distance $\Delta\vartheta$ in that region, and so is the periodicity. Locally, the smoothing function has a sinc-like shape, similarly to the case of the linear pole distribution of Sec. 5.1. The dashed lines show the theoretical $1/|\vartheta - \vartheta_0|$ envelope of the sinc function. Again, (b) is a borderline case where the envelope still follows that of a regular sinc function, but the periodicity is different at the left and right side, coming from the different pole densities.

5.3. Logarithmic pole distribution

As a final example, let us consider a case with logarithmic frequency resolution, when the parallel filter has three poles in each octave, having all together 31 pole pairs from 20 to 20480 Hz. The test function is again a complex exponential, with $\vartheta_0 = 2\pi f_0/f_\mathrm{s}$, where $f_0 = 1050$ Hz is the frequency of the exponential, and $f_\mathrm{s} = 44.1$ kHz is the sampling frequency. The time domain (real part) and frequency domain responses are displayed in Fig. 5 (a) and (b), respectively. Note the linear frequency axis in (b).

It can be seen in Fig. 5 (a) that now the target function is "windowed" quite mildly and it has a low frequency tail. In the frequency domain (Fig. 5 (b)) the density of the notches of the *sinc*-like function (solid line) follow that of the pole distribution (dotted vertical lines). However, the envelope of the smoothing function still shows the $1/\vartheta$ behaviour, corresponding to the envelope of a regular sinc function.

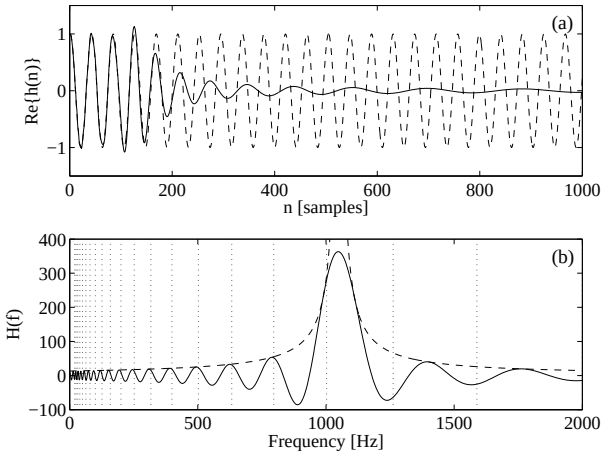

Fig. 5: Modeling an exponential test function $e^{-j\vartheta_0 n}$ with a parallel filter having logarithmic frequency resolution: (a): the real parts of the target impulse response (dashed line) and the parallel filter response (solid line), (b): the smoothing function of the parallel filter (solid line) and its $1/|\vartheta - \vartheta_0|$ envelope. The vertical dotted lines display the pole frequencies.

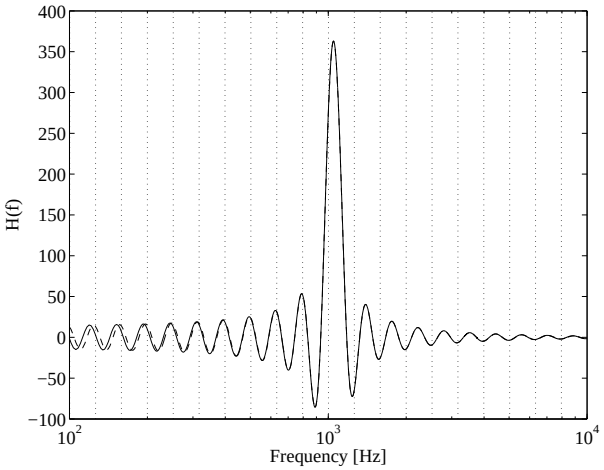

Fig. 6: The smoothing function of Fig. 5 (b) displayed in a logarithmic frequency axis (solid line). The dashed line shows the approximating "*logsinc*" function Eq. (21). The vertical dotted lines display the pole frequencies.

Figure 6 solid line displayes the same frequency response on a logarithmic frequency scale. Now it is easy to note that the periodicity of the window function is exactly logarithmic. The smoothing function may be approximated by a "*logsinc*" function

$$S(\vartheta) = C\frac{\sin(2\pi\beta\log_2(\frac{\vartheta}{\vartheta_0}))}{\vartheta - \vartheta_0}, \qquad (21)$$

where C is a positive constant, and β is the pole per octave density (in our case, $\beta = 3$). This function is displayed by a dashed line in Fig. 6, matching the filter response very precisely.

Naturally, further examples could be presented with, e.g., stepwise logarithmic pole distribution, or that of following the Bark or ERB scale, but according to the above examples, the reader should already have an intuition about the smoothing behavior of the parallel filter.

6. DESIGN EXAMPLES AND COMPARISON

An example for loudspeaker-room response modeling has already been displayed in Fig. 1. Now we are designing an equalizer for the same system response. The target response is a second-order high-pass filter with a cutoff frequency of 50 Hz, and the frequency resolution of the design is third octave.

The system response $h_{\mathrm{s}}(n)$ is displayed in Fig. 7 (a), 1/12 octave smoothed for clarity. It's third-octave complex smoothed version $h_{\mathrm{cs}}(n)$ is displayed in Fig. 7 (a), which is used for a 2500 tap FIR equalizer design by a least-squares system identification approach [12]. In this case, the parameters are estimated in such a way so that the complex-smoothed system response $h_{\mathrm{cs}}(n)$ filtered by the FIR filter best approximates the target response $h_{\mathrm{t}}(n)$ in the LS sense. The system response equalized by this FIR filter is shown in Fig. 7 (b), while the transfer function of the FIR equalizer itself can be seen in Fig. 8 (b). The FIR filter smooths the overall response quite nicely while avoids equalizing the narrow peaks and dips that would happen when designing from the unsmoothed system response (direct inversion).

However, the computational complexity of the 2500 tap FIR filter is too large compared to the simple task it should accomplish. A straightforward option for decreasing the computational complexity is to estimate a warped IIR filter based on the complex-smoothed system response. Figure 7 (c) displays the loudspeaker-room response equalized by such a WIIR filter having an order of 32 and $\lambda = 0.75$. The transfer function of the WIIR filter is shown in Fig. 8 (c). It can be seen in Figs. 7 and 8 that

the results are almost the same as for the FIR filter, expect at very low frequencies. Note that the low frequency precision could be increased by using larger λ values, but that often results in unstable filters in practice. Also note that the case of strictly logarithmic frequency resolution coming from third-octave smoothing is a relatively easy task for the WIIR filter, since it is quite close to its natural frequency resolution. For arbitrary frequency resolution profiles (such as having much higher resolution at low frequencies compared to the high ones), the performance of the WIIR filter would be less than satisfactory.

Figure 7 (d) shows the loudspeaker-room response equalized by a 16 section parallel filter, having the same total filter order as the WIIR filter of (c). The parallel filter was designed by the direct method presented in Sec. 4.3, by using the *unsmoothed* system response (we note again that the FIR and WIIR examples were designed using the *complex-smoothed* system response). To achieve the required third-octave resolution, three poles are placed in each two octaves (the pole density is 3/2 pole/octave). The resulting equalizer performance is as good as for the FIR and IIR case, while the design is much simplified, since there is no need for smoothing the target response. This is because the smoothing is done "automatically", as already discussed in Sec. 2. In addition, the procedure always results in a stable filter, and is equally efficient at arbitrary (i.e., not strictry logaritmic) frequency resolution, on the contrary to the WIIR design.

Figure 9 shows two additional examples illustrating some of the capabilities of the parallel filter using arbitrary pole distributions. Often it is desired that the low frequencies of the loudspeaker-room response are more precisely equalized compared to the high ones. This is now achieved by a parallel filter having 6 poles per octave below 320 Hz (corresponding to twelfth-octave resolution), and 3/2 pole per octave above (third-octave resolution). The equalized response is displayed in Fig. 9 (b), while the transfer function of the equalizer is shown in (d), together with the pole frequencies as vertical dotted lines. One can see that the 34 section parallel filter achieves a quite flat frequency response.

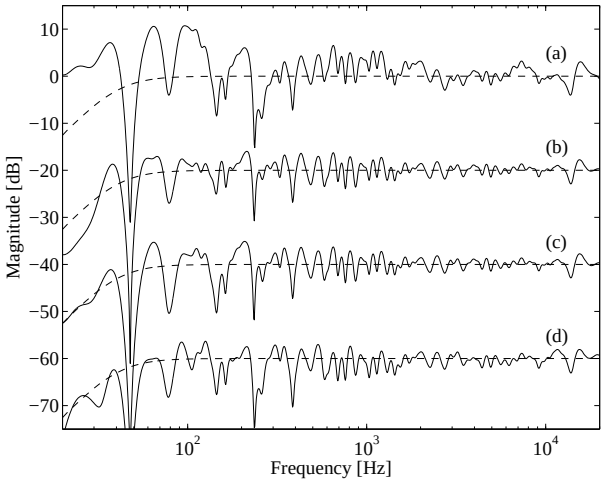

Fig. 7: Loudspeaker-room equalization: (a) the unequalized system response, (b) equalized by a 2500 tap FIR filter estimated using the third-octave complex-smoothed system response, (c) 32nd order WIIR filter estimated using the third-octave complex-smoothed system response, and (d) a 16 section (32nd order) parallel filter designed using the *unsmoothed* system response. All responses smoothed to a 1/12 octave resolution for clarity.

Figure 9 (c) displays an example when only the problematic low frequency region is equalized by seven second-order sections. In this case, the seven pole pairs are logarithmically distributed between 20 and 320 Hz, corresponding to a third-octave resolution (see dotted vertical lines in Figure 9 (e)), and a zero-order FIR part (the constant coefficient b_0 in Fig. 2) is also utilized. The transfer function of the equalizer is shown in Fig. 9. It is important to note that the equalizer is designed from the unsmoothed loudspeaker-room response exactly as in the previous cases, and there is no need to do any additional processing, like flattening the response above the frequency region of the equalization, etc.

7. DISCUSSION

We have seen in Sec. 2 that designing a parallel filter is equivalent to smoothing the target response by a sinc-like function, and in the examples of Sec. 6 that the smoothing behavior is also present in direct equalizer

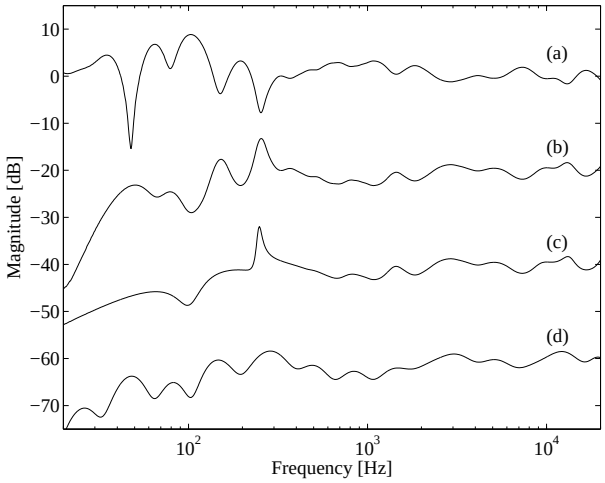

Fig. 8: Equalizer responses for the loudspeaker-room equalization of Fig. 7: (a) the third-octave smoothed system response used for designing the FIR and WIIR equalizers, the transfer functions of the (b) FIR and (c) WIIR equalizers, (d) and the transfer function of the parallel filter.

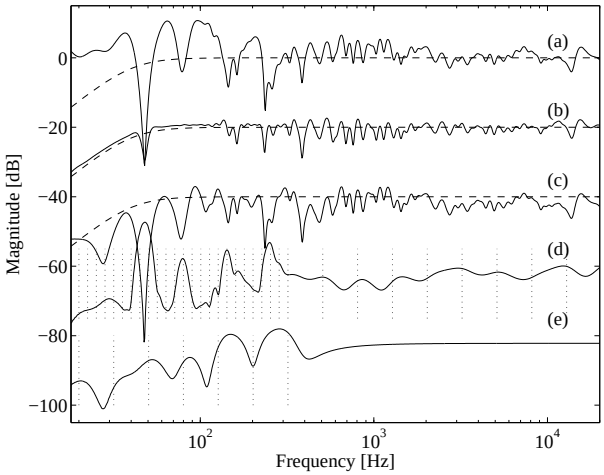

Fig. 9: Illustrating the possibilities using various pole distributions in the parallel filter: (a) the unequalized loudspeaker-room response, (b) equalized by a parallel filter having four times higher pole density below 320 Hz then above, (c) equalized by a parallel filter having poles only below 320 Hz. The transfer functions of the equalizers for the cases (b) and (c) are displayed in (d) and (e), respectively. The dotted vertical lines display the pole frequencies of the corresponding equalizers. Curves (a)–(c) are smoothed to a 1/12 octave resolution for clarity.

design. The width of the main lobe of the sinc function equals to the pole frequency distance in that region. Moreover, the main lobe of the sinc function closely approximates a hanning window, a window function commonly used in transfer function smoothing. For obtaining a given Δf resolution at some frequency region, the distance of the analog pole frequencies of the parallel filter should be $2\Delta f$, since that corresponds to smoothing by a $2\Delta f$ wide hanning window. In the logarithmic frequency scale, if $1/\beta$ octave resolution is desired, $\beta/2$ pole pairs have to be placed in each octave.

In practice, the fixed-pole parallel filters can be used in two ways in relation to complex smoothing:

1) Design of parallel filters instead of transfer function smoothing: In this case, the system is modeled or equalized by the parallel filter without any frequency-domain processing. Thus, the parallel filter is used both as the final implementation structure and a means of achieving complex smoothing. Here, the resolution is controlled by the choice of the pole frequencies. The filter is optimal in the sense that the filter order will correspond to the obtained resolution. An additonal advantage of this approach that the parameter estimation is simplified since there is no need to implement the complex-smoothing operation.

2) Design of parallel filters together with transfer function smoothing: In this case, the parallel filter is designed by using the smoothed frequency response. This does not make too much sense for complex-smoothed responses, since almost the same results could be achieved without prior smoothing. However, if the system response is smoothed by some nonlinear processing (e.g., eliminating the dips below certain threshold, or by the use of an auditory model), then the parallel filter can be used as an efficient implementation of the smoothed response. In this case, the local pole density of the parallel filter should be set according to the local resolution of the smoothed transfer function. (Note that for implementing already smoothed responses the warped IIR filter is also a strong contender, but it may result in unstable filters, and may provide unsatisfactory results for such smoothing profiles that are not close to the natural frequency resolution of the warped filter.)

8. CONCLUSION

Transfer function smoothing is a well-established method for displaying, modeling and equalizing the frequency responses of audio systems, coming from both perceptual and physical considerations. This paper has demonstrated that the impulse response of the parallel filter is similar to the response obtained by the complex smoothing of the target response. As a result, the parallel filter can be either used as an efficient implementation of already smoothed responses, or, it can be designed from the unsmoothed responses directly, eliminating the need of frequency-domain processing, since it performs smoothing "automatically".

The obtained frequency resolution is not limited to the logarithmic scale, but arbitrary resolution can be achieved by the suitable choice of pole frequencies. The formulas for computing the pole angles and radii from analog pole frequencies were also given.

The theoretical equivalence of parallel filters and Kautz filters has also been developed, and the formulas for converting between the parameters of the two structures were presented. This implies that the favorable smoothing properties are also possessed by the Kautz filter. In addition, the conversion formulas can be used for obtaining the parameters of the parallel filter from the Kautz parameters, resulting in a design procedure that requires less arithmetic operations compared to the straightforward LS design.

While only loudspeaker-room equalization examples have been provided, the parallel filter can be successfully used also in other fields where the flexible allocation of frequency resolution is beneficial. So far, it has been applied to modeling the body radiation [19] and bridge admittance [22] of musical instruments for sound synthesis, and the modeling of the direction dependent radiation of guitar speakers [23], and it is hoped that other applications will soon follow.

9. ACKNOWLEDGEMENT

This work has been supported by the EEA Norway Grants and the Zoltán Magyary Higher Education Foundation.

10. REFERENCES

[1] J. N. Mourjopoulos, P. M. Clarkson, and J. K. Hammond. A comparative study of least-squares and homomorphic techniques for the inversion of mixed phase signals. In *Proc. IEEE Int. Conf. Acoust. Speech and Signal Process.*, pages 1858–1861, May 1982.

[2] Matti Karjalainen, Esa Piirilä, Antti Järvinen, and Jyri Huopaniemi. Comparison of loudspeaker equalization methods based on DSP techniques. *J. Audio Eng. Soc.*, 47(1–2):14–31, Jan./Feb. 1999.

[3] German Ramos and Jose J. Lopez. Filter design method for loudspeaker equalization based on IIR parametric filters. *J. Audio Eng. Soc.*, 54(12):1162–1178, Dec. 2006.

[4] Masato Miyoshi and Yutaka Kaneda. Inverse filtering of room acoustics. *IEEE Trans. Acoust. Speech Signal Process.*, 36(2):145–152, Feb. 1988.

[5] Peter G. Craven and Michael A. Gerzon. Practical adaptive room and loudspeaker equalizer for hi-fi use. In *Proc. 92nd AES Conv., Preprint No. 3346*, Vienna, Austria, Mar. 1992.

[6] 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.

[7] Sunil Bharitkar and Chris Kyriakakis. Loudspeaker and room response modeling with psychoacoustic warping, linear prediction, and parametric filters. In *Proc. 121st AES Conv., Preprint No. 6982*, San Francisco, CA, USA, Oct. 2006.

[8] Jan Abildgaard Pedersen and Kasper Thomsen. Fully automatic loudspeaker–room adaptation – The RoomPerfect system. In *Proc. AES 32nd Int. Conf. "DSP for Loudspeakers"*, pages 11–20, Hillerød, Denmark, Sep. 2007.

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

[10] Matti Karjalainen and Tuomoas Paatero. Frequency-dependent signal windowing. In *Proc. IEEE Workshop Appl. of Signal Process. to Audio and Acoust.*, pages 35–38, New Paltz, NY, USA, Oct. 2001.

[11] Joerg Panzer and Lampos Ferekidis. The use of continuous phase for interpolation, smoothing and forming mean values of complex frequency response curves. In *Proc. 116th AES Conv., Preprint No. 6005*, Berlin, Germany, 2004.

[12] J. N. Mourjopoulos and P. D. Hatziantoniou. Real-time room equalization based on complex smoothing: Robustness results. In *Proc. 116th AES Conv., Preprint No. 6070*, May 2004.

[13] Michael Waters and Mark B. Sandler. Least squares IIR filter design on a logarithmic frequency scale. In *Proc. IEEE Int. Symp. on Circuits and Syst.*, pages 635–638, May 1993.

[14] 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.

[15] Julius O. Smith and Jonathan S. Abel. Bark and ERB bilinear transform. *IEEE Trans. Speech Audio Process.*, 7(6):697–708, Nov. 1999.

[16] Stanley P. Lipshitz, Tony C. Scott, and John Vanderkooy. Increasing the audio measurement capability of fft analyzers by microprocessor postprocessing. *Journal of Applied Mechanics*, 33(9):626–648, Sep. 1985.

[17] 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.

[18] Matti Karjalainen and Tuomas Paatero. Equalization of loudspeaker and room responses using Kautz filters: Direct least squares design. *EURASIP J. on Advances in Sign. Proc., Spec. Iss. on Spatial Sound and Virtual Acoustics*, 2007:13, 2007. Article ID 60949, doi:10.1155/2007/60949.

[19] 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`.

[20] 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`.

[21] L. R. Rabiner and B. Gold. *Theory and Application of Digital Signal Processing*. Prentice-Hall, Englewood Cliffs, New Jersey, USA, 1975.

[22] Balázs Bank and Matti Karjalainen. Passive admittance synthesis for sound synthesis applications. In *Proc. Acoustics'08 Paris Conf.*, Paris, France, June 2008. paper id: ACOUSTICS2008/254.

[23] David Yeh, Balázs Bank, and Matti Karjalainen. Nonlinear modeling of a guitar loudspeaker cabinet. In *Proc. Conf. on Digital Audio Effects*, pages 89–96, Espoo, Finland, Sep. 2008.