Audio Engineering Society

Convention Paper

Presented at the 136th Convention
2014 April 26–29 Berlin, Germany

This Convention paper was selected based on a submitted abstract and 750-word precis that have been peer reviewed by at least two qualified anonymous reviewers. The complete manuscript was not peer reviewed. This convention paper has been reproduced from the author's advance manuscript without editing, corrections, or consideration by the Review Board. The AES takes no responsibility for the contents. Additional papers may be obtained by sending request and remittance to Audio Engineering Society, 60 East 42nd Street, New York, New York 10165-2520, USA; also see www.aes.org. All rights reserved. Reproduction of this paper, or any portion thereof, is not permitted without direct permission from the Journal of the Audio Engineering Society.

A delayed parallel filter structure with an FIR part having improved numerical properties

Balázs Bank[1] and Julius O. Smith[2]

[1] *Budapest University of Technology and Economics, Department of Measurement and Information Systems, H-1521 Budapest, Hungary*

[2] *Center for Computer Research in Music and Acoustics, Stanford University, Stanford, CA 94305-8180, USA*

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

ABSTRACT

In real-world applications high-order IIR filters are often converted to series or parallel second-order sections to decrease the negative effects of coefficient truncation and roundoff noise. While series biquads are more common, the parallel structure is gaining more interest due to the possibility of full code parallelization. In addition, it is relatively simple to design a filter directly in a parallel form, which can be efficiently utilized for logarithmic frequency resolution filtering often required in audio. If the numerator order of the original transfer function is higher than that of the denominator, a parallel FIR part arises in addition to the second-order IIR sections. Unfortunately, in this case the gain of the sections and that of the FIR filter can be significantly higher than that of the final transfer function, which requires the downscaling of the filter coefficients to avoid overload. This leads to a significant loss of useful bitdeph. This paper analyzes problem and suggests delaying the IIR part so that there is no overlap between the responses of the FIR part and the second-order sections.

1. INTRODUCTION

High-order IIR filters are used in many fields of audio, including the modeling and equalization of loudspeaker and room responses, HRTF modeling, sound synthesis, etc. However, in real-world applications where double-precision floating point arithmetic is usually not available, direct form high-order IIR filters lead to a poor numerical performance and often become unstable due to coefficient rounding effects. Therefore, they are usually converted to series or parallel second-order sections to decrease the negative effects of coefficient rounding and roundoff noise [1, 2, 3, 4]. While series biquads are more common, the parallel structure is gaining more and more interest due to the possibility of full code parallelization.

The parallel form is usually computed by the help of partial fraction expansion [1, 2], but it is also possible to use an iterative procedure for conversion [5]. In either case, if the numerator order of the original transfer function is higher than that of the denominator, a parallel FIR part arises besides the second-order IIR sections.

In addition, parallel second-order filters can be designed directly in a closed form if the pole positions are predetermined [6]. This kind of design methodology is connected to warped [7] and Kautz [8] filters where the goal is to obtain a filter with a logarithmic frequency resolution, better suited to the properties of hearing than traditional, linear frequency resolution filters. It was shown in [9] that the frequency resolution of the fixed-pole parallel filter is directly controlled by the pole positions. For example, fixing the poles according to a logarithmic frequency scale results in a logarithmic frequency resolution. In the context of fixed-pole parallel filters, a parallel FIR part is used when the target is largely non-minimumphase, such as modeling the radiation transfer function of musical instruments [6], or in non-minimumphase loudspeaker equalization.

However, either converted from a high-order direct-form IIR filter or designed by the fixed-pole method, a dynamic range problem can arise with parallel filters if there is an FIR part present. In this case the gain of the second-order sections and that of the FIR filter can be significantly higher than that of the final transfer function, which requires the downscaling of the filter coefficients to avoid overload. This leads to a significant loss of useful bitdeph.

The rest of the paper is organized as follows: Sec. 2 describes the usual procedure for obtaining second-order parallel filters from high-order direct-form filters, Sec. 3 summarizes the fixed-pole design of parallel filters, and Sec. 4 analyzes the dynamic range problem when there is an FIR part. Section 5 proposes using a delayed parallel form and discusses parameter estimation for that structure, and finally Sec. 6 compares the performance of the straightforward and the proposed delayed parallel filters.

2. PARALLEL FILTERS BY PARTIAL FRACTION EXPANSION

The usual way of converting direct form IIR filters to parallel sections is by the use of *partial fraction expansion* [1, 2, 4].

Suppose we start with the direct-form transfer function

$$H(z) \triangleq \frac{B(z)}{A(z)} = \sum_{i=1}^{N} \frac{r_i}{1-p_i z^{-1}} \qquad (1)$$

where

$$\begin{aligned} B(z) &= b_0 + b_1 z^{-1} + b_2 z^{-2} + \cdots + b_M z^{-M} \\ A(z) &= 1 + a_1 z^{-1} + a_2 z^{-2} + \cdots + a_N z^{-N}, \end{aligned}$$

and suppose initially that $M < N$. The denominator coefficients p_i are the *poles* of the transfer function, and each numerator r_i is called the *residue* of pole p_i. Equation (1) is general only if the poles p_i are *distinct*. (Repeated poles are addressed in [4].) The poles and their residues may be complex. The poles p_i may be found by factoring the polynomial $A(z)$ into first-order terms. The residue r_i corresponding to pole p_i is given by

$$r_i = (1-p_i z^{-1})H(z)\big|_{z=p_i} \qquad (2)$$

when the poles p_i are distinct. The function `residuez` in Octave or the Matlab Signal Processing Toolbox will find poles and residues computationally, given the direct-form coefficients $\{b_i, a_i\}$.

2.1. FIR Part

When $M \geq N$ in Eq. (1), we may perform a step of *long division* of $B(z)/A(z)$ to produce an *FIR part* in parallel with a strictly proper IIR part:

$$H(z) \triangleq \frac{B(z)}{A(z)} = F(z) + \sum_{i=1}^{N} \frac{r_i}{1-p_i z^{-1}} \qquad (3)$$

where

$$\begin{aligned} B(z) &= b_0 + b_1 z^{-1} + b_2 z^{-2} + \cdots + b_M z^{-M} \\ A(z) &= 1 + a_1 z^{-1} + a_2 z^{-2} + \cdots + a_N z^{-N} \\ F(z) &= f_0 + f_1 z^{-1} + f_2 z^{-2} + \cdots + f_K z^{-K}, \end{aligned}$$

where $K = M - N$. When $M < N$, we define $F(z) = 0$. This type of decomposition is computed by the `residuez` function.

In summary, an arbitrary digital filter transfer function $H(z)$ with N distinct poles can always be expressed as a parallel combination of *complex one-pole filters*, together with a parallel FIR part when $M \geq N$.

2.2. Second-order sections with real coefficients

Instead of implementing complex first-order sections it is more practical to combine the complex-conjugate pole pairs to second-order sections, so the transfer function becomes

$$H(z) = \sum_{l=1}^{L} \frac{b_{0,l} + b_{1,l} z^{-1}}{1 + a_{1,l} z^{-1} + a_{2,l} z^{-2}} + \sum_{k=0}^{K} f_k z^{-k}, \quad (4)$$

where L is the number of sections. For complex pole pairs $p_i = p_l$ and $p_{i+1} = \overline{p}_l$ the second-order coefficients are obtained as follows:

$$b_{0,l} = 2\mathrm{Re}\{r_l\} \quad (5)$$
$$b_{1,l} = -2\mathrm{Re}\{r_l \overline{p}_l\} \quad (6)$$
$$a_{1,l} = -2\mathrm{Re}\{p_l\} \quad (7)$$
$$a_{2,l} = |p_l|^2 \quad (8)$$

Sections having real poles in Eq. (1) can be directly implemented as first-order IIR filters, but it is customary to combine them in pairs for a coherent filter structure containing only second-order sections plus at most one first-order section.

3. PARALLEL FILTERS BY FIXED-POLE DESIGN

In fixed-pole design of parallel filters, the filter coefficients are estimated directly in the parallel form as in Eq. (4) for a predefined set of poles. The advantage of fixed-pole design is that by the choice of pole frequencies we gain complete control over the frequency resolution of the filter. For example, placing the poles according to a logarithmic frequency grid results in a filter with a logarithmic frequency resolution, similar to fractional octave smoothing, a procedure often used in audio for displaying, modeling, and equalizing acoustic transfer functions [10, 9]. The discussion of various pole positioning methods for the parallel filter is out of scope of this paper, and a thorough comparison can be found in [11].

Once the denominator coefficients are determined by the poles ($a_{l,1} = -2\mathrm{Re}\{p_l\}$ and $a_{l,2} = |p_l|^2$), the problem becomes linear in its free parameters $b_{l,0}$, $b_{l,1}$ and f_k in Eq. (4). The filter can be designed either based on a target impulse response [12] or a target frequency response [13]. Here the former will be outlined.

The impulse response of the parallel filter is given by

$$h(n) = \sum_{l=1}^{L} b_{l,0} u_l(n) + b_{l,1} u_l(n-1) + \sum_{k=0}^{K} f_k \delta(n-k) \quad (9)$$

where $u_l(n)$ is the impulse response of the transfer function $1/(1 + a_{l,1} z^{-1} + a_{l,2} z^{-2})$, which is an exponentially decaying sinusoidal function or pair of exponential decays, and $\delta(n)$ is the discrete-time unit impulse.

Naturally, Eq. (9) is linear in the parameters, similar to its z-transform counterpart Eq. (4). Writing Eq. (9) in matrix form yields

$$\mathbf{h} = \mathbf{M}\mathbf{p} \quad (10)$$

where $\mathbf{p} = [b_{1,0}, b_{1,1}, \ldots b_{L,0}, b_{L,1},\ f_0 \ldots f_K]^T$ is a column vector composed of the free parameters. The columns of the modeling signal matrix $\mathbf{M}$ contain the modeling signals, which are $u_l(n)$ and their delayed counterparts $u_l(n-1)$, and for the FIR part, the unit impulse $\delta(n)$ and its delayed versions up to $\delta(n-K)$. 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 mean squared error is minimized, the optimum is found by the well known LS solution

$$\mathbf{p}_{\mathrm{opt}} = (\mathbf{M}^H \mathbf{M})^{-1} \mathbf{M}^H \mathbf{h}_{\mathrm{t}} \quad (11)$$

where $\mathbf{M}^H$ is the conjugate transpose of $\mathbf{M}$.

4. THE HIGH GAIN PROBLEM

Either converted from a high-order direct-form IIR filter or designed by the fixed-pole method, a dynamic range problem can arise with parallel filters if there is an FIR part present. This is illustrated in Fig. 1, showing a loudspeaker response modeling example. First, an IIR filter is designed with $M = 50$ and $N = 20$ by the Steiglitz-McBride method [14] (`stmcb` function in the Matlab Signal Processing Toolbox), then the filter is converted to parallel form Eq. (4) by partial fraction expansion as described in Sec. 2. The specification (thin solid line) is followed by the filter response (thick line) quite accurately. However, when plotting the responses of the individual second-order IIR sections (dotted lines) and the 30th order FIR filter (dashed line) separately, it can

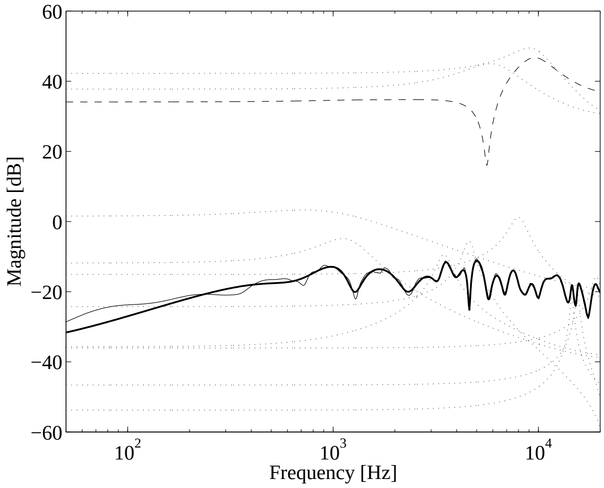

Fig. 1: Parallel filter obtained by partial fraction expansion for modeling a loudspeaker response. Thin line: target frequency response, thick line: filter response, dashed line: frequency response of the FIR part, dotted lines: the individual responses of the second-order sections.

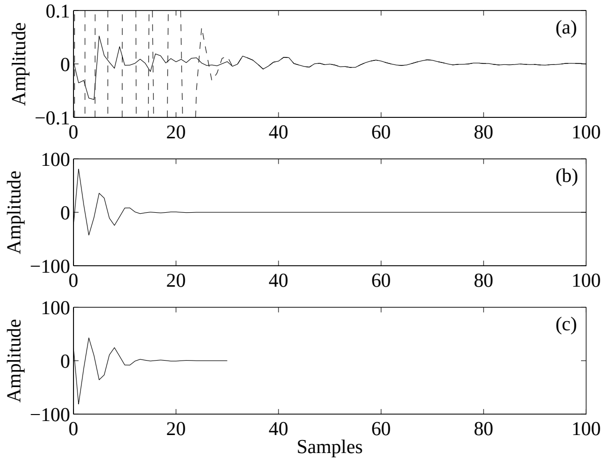

Fig. 2: Time-domain responses of the parallel filter used in Fig. 1: (a) filter impulse response (solid line) and the summed impulse response of the second-order sections (dashed line, clipping), (b) the response of the second-order sections in a different scale, and (c) the response of the 30th order FIR part.

be seen that the gain of the sections can easily exceed the gain of the total transfer function by 60 dB, and thus the total transfer function is formed by cancellations of the different parts. Intuitively, the cancellation can be understood by considering that the FIR part will match the early impulse-response (IR) exactly, while the IIR part is only needed to approximate the later IR, beyond the FIR duration. Therefore, any decay of the IIR part during the FIR part must be canceled by the FIR part. This is illustrated in the time-domain by Fig. 2 for the same filter as in Fig. 1. It can be seen in Fig. 2 (a) that the IIR part (dashed line) is much larger than the final impulse response (solid line), and that the FIR part (c) is basically the inverse of the IIR part (b).

The worst combination is a long FIR part followed by a highly damped IIR “tail”; in such a case, a very large initial IIR part must be canceled by the FIR part. This will be the case for the piano soundboard modeling example in Sec. 6. In the opposite direction, as the poles of IIR part approach magnitude 1, or as the FIR part is made shorter, there is less decay across the FIR duration, alleviating the dynamic range problem.

This FIR-IIR cancellation leads to two distinct problems in practical applications: first, the filter coefficients (or the input signal, or both) need to be downscaled so that no overload occurs in the separate transfer functions, and second, the total transfer function will be more sensitive to the coefficient quantization errors since it arises as the difference of large, close-to-each-other values. Note that while the downscaling is not necessary in floating point arithmetic, the loss of useful bitdepth still occurs. In this particular case, the 60 dB difference leads to losing as many as 10 bits.

The reason for this behavior is more easily understood when examined from the fixed-pole design point of view. In the fixed-pole design of parallel filters a linear least squares (LS) fit is performed where the basis functions are delayed impulses for the FIR part and exponentially decaying sinusoids (and their one-sample delayed versions) for the second-order IIR sections (see Sec. 3). The weights are set so that the filter impulse response best approximates the target impulse response. Since a Kth order FIR part gives complete freedom for setting the first $K+1$ samples of the filter response, the LS design will set the FIR coefficients in such a way that the first $K+1$ samples ($n=[0..K]$) are matched perfectly. This means that the numerators of the second-order sections will depend only on the samples after the FIR part ($n > K+1$). However, if the FIR part is sufficiently long, the decaying sinusoids have already a low level at this sample point, which can be counteracted by the LS design only by increasing the initial amplitudes of the sinusoids, and thus the numerator coefficients of the second-order filters. In return, this will mean a large signal at the beginning of the response, overlapping the FIR part. Indeed, besides setting the initial sample values, the main role of the FIR coefficients is to cancel the response of the IIR part in the first $K+1$ samples, which will mean that the FIR coefficients will also be much larger than the target impulse response.

This least squares procedure can also be used for converting a high-order direct-form IIR filter to a parallel second-order form once the poles of the direct-form IIR filter are computed, and since the solution is unique and has no approximation error, it will be the same as obtained by partial fraction expansion (if we neglect numerical errors). Therefore, the above reasoning is also valid when the parallel filter is obtained by partial fraction expansion.

5. THE DELAYED PARALLEL FILTER

Once understood, this problem can be solved in a simple way: the parallel IIR part must be delayed so that there is no overlap between the FIR and IIR parts:

$$H(z) = z^{-(K+1)} \sum_{l=1}^{L} \frac{\tilde{b}_{l,0} + \tilde{b}_{l,1} z^{-1}}{1 + a_{l,1} z^{-1} + a_{l,2} z^{-2}} + \sum_{k=0}^{K} \tilde{f}_k z^{-k}, \tag{12}$$

The first $K+1$ samples of the impulse response are now determined solely by the the Kth order FIR part, and the rest of the impulse response by the IIR part.

Note that this delayed form does not require additional memory since the last element of the delay line used for the FIR part can be used as an input for the parallel second-order IIR part. More precisely, one extra delay element is needed, since the largest term in the FIR part is z^{-K} and we need $z^{-(K+1)}$.

In the following we discuss how the parameters of the delayed parallel filter can be obtained.

5.1. Obtaining filter coefficients from the non-delayed parallel filter

The parameters of the delayed filter structure can be obtained from the original parallel structure of Eq. (4) with a Kth order FIR part as follows: the first $K+1$ samples of the filter impulse response $h(k)$ are computed, and these samples are directly used as the new FIR coefficients:

$$\tilde{f}_k = h(k) \text{ for } k = [0, 1, \ldots K+1] \tag{13}$$

For the parallel IIR sections, the denominators remain the same and the numerators are set in such a way that the decaying exponentials have the same amplitude and phase at sample $n = 0$ as at sample $n = K+1$ with the original sections. First the second-order sections are converted to a pair of complex first-order IIR filters as in Eq. (1). Then the modified residues are obtained as

$$\tilde{r}_i = r_i p^{K+1}, \tag{14}$$

and finally the first-order filters are again combined to form second-order sections having real coefficients.

5.2. Alternative way to convert from direct-form IIR filters

The delayed parallel filter can also be computed directly by modified partial fraction expansion. The alternate FIR part is obtained by performing long division on the *reversed* polynomial coefficients to get

$$H(z) = F(z) + z^{-(K+1)} \sum_{i=1}^{N} \frac{r_i}{1 - p_i z^{-1}},$$

where $K = M - N \geq 0$ is again the order of the FIR part [4].

We may compare these two PFE alternatives as follows: Let A_N denote $A(z)$, $F_K \triangleq F(z)$, and $B_M \triangleq B(z)$. (*I.e.*, we use a subscript to indicate polynomial order, and '(z)' is omitted for notational simplicity.) Then for $K = M - N \geq 0$ we have two cases:

$$(1) \quad H(z) = F_K + \frac{B'_{N-1}}{A_N} = \frac{F_K A_N + B'_{N-1}}{A_N}$$

$$(2) \quad H(z) = F_K + z^{-(K+1)} \frac{B''_{N-1}}{A_N} = \frac{F_K A_N + z^{-(K+1)} B''_{N-1}}{A_N}$$

In the first form, the B'_{N-1} coefficients are "left justified" in the reconstructed numerator, while in the second form they are "right justified". In addition to avoiding the high-gain problem of the first form, the second form is generally more efficient for *modeling* purposes, since the numerator of the IIR part ($B''_{N-1}(z)$) can be used to match additional terms in the impulse response after the FIR part $F_K(z)$ has "died out".

Method

Figure 3 gives a listing of the matlab function `residued` (distributed with Octave) for computing a "right justified" partial fraction expansion (PFE) of an IIR digital filter $H(z) = B(z)/A(z)$ as described in Sec. 2.

```
function
  [r, p, f, e] = residued(b, a, toler)
if nargin<3, toler=0.001; end
NUM = b(:)';
DEN = a(:)';
nb = length(NUM);
na = length(DEN);
f = [];
if na<=nb
 f = filter(NUM,DEN,[1,zeros(nb-na)]);
 NUM = NUM - conv(DEN,f);
 NUM = NUM(nb-na+2:end);
end
[r,p,f2,e] = residuez(NUM,DEN,toler);
```

Fig. 3: Matlab/Octave function for computing a partial fraction expansion in which the parallel sections follow the FIR part.

As can be seen in Figure 3, the FIR part is first extracted, and the (strictly proper) remainder is passed to `residuez` for expansion of the IIR part (into a sum of complex resonators).

Alternatively, a delayed partial fraction form can also be obtained by computing the partial fraction expansion with `residue` in Matlab/Octave, which is intended for analog (s domain) transfer functions. We note however that when there is no FIR part ($M < N$) this produces wrong r_i residues and for that `residuez` has to be used. An advantage of `residued` in Figure 3 is that for $M \geq N$ it gives the parameters for the delayed parallel filter, but it is also correct for the strictly proper (FIR-less) case ($M < N$).

5.3. Modified fixed-pole design of parallel filters

When using fixed-pole parallel filters it is also possible to design the filter in the delayed form of Eq. (12) directly. In this case, we choose the FIR coefficients $\tilde{f}_k$ as the first $K + 1$ samples of the target impulse response $h_\mathrm{t}(k)$:

$$\tilde{f}_k = h_\mathrm{t}(k) \text{ for } k = [0, 1, \ldots K + 1], \qquad (15)$$

where K is the order of the FIR part.

Then, the remaining part of the target will be used as a specification

$$\tilde{h}_\mathrm{t}(n) = h_\mathrm{t}(n + K + 1) \qquad (16)$$

for designing a strictly proper ($M < N$) parallel filter. Thus, the modeling matrix $\tilde{\mathbf{M}}$ contains only the responses corresponding to the second-order sections, but no unit pulses.

Besides avoiding the conversion, this has an added benefit of decreasing the computational complexity of the design, since now that the normal equations estimate $K+1$ fewer parameters. In addition, the design problem becomes numerically better conditioned, since in the original design of Sec. 3 the same samples were determined by the FIR and the IIR parts, which is now avoided.

Note that when the delayed parallel filter is designed based on a frequency response specification, we cannot take advantage of the fact that the FIR coefficients are the same as the early part of the impulse response. In this case all free parameters (FIR and numerator coefficients) have to be computed by the LS solution, similarly as done in [13]. The only difference is that the frequency responses of the second-order sections in the frequency-domain modeling matrix $\mathbf{M}$ used in [13] are multiplied by $z^{-(K+1)} = e^{-j\vartheta(K+1)}$ in accordance with Eq. (12).

6. COMPARISON

6.1. Partial fraction example

When the same filter as in Fig. 1 is converted to a parallel filter with a delayed IIR part by the alternative form of partial fraction expansion as in Sec. 5.2, the high gain problem is avoided. This is displayed in Fig. 4. It can be seen that none of the individual transfer functions exceed the total filter response.

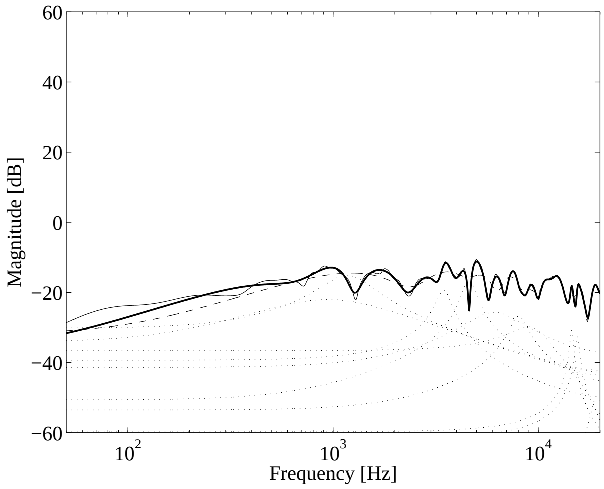

Fig. 4: Delayed parallel filter obtained by the alternative partial fraction expansion of Sec. 5.2 from the same direct-form filter as in Fig. 1. Thin line: target frequency response, thick line: filter response, dashed line: frequency response of the FIR part, dotted lines: the responses of the second-order sections.

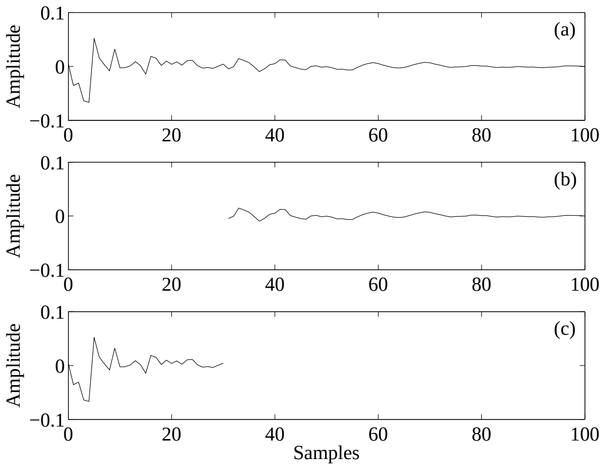

Fig. 5: Time-domain responses of the parallel filter used in Fig. 4: (a) filter impulse response (solid line), (b) the response of the second-order sections, and (c) the response of the 30th order FIR part.

The same responses are displayed in the time-domain in Fig. 5, showing how the delayed IIR (b) and FIR (c) parts are combined to form the total impulse response (a).

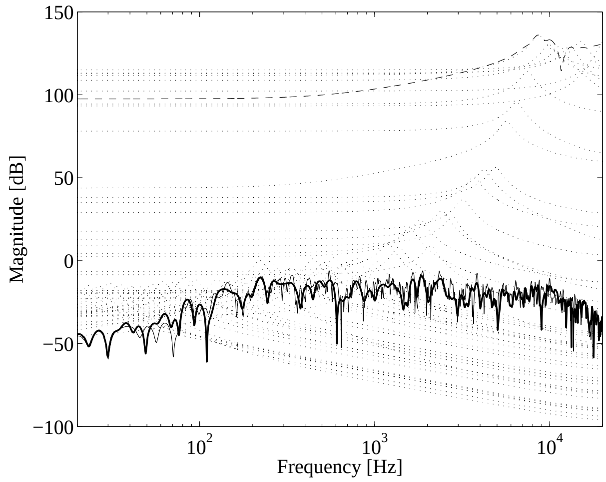

Fig. 6: Fixed-pole parallel filter design based on a piano soundboard response. Thin line: target frequency response, thick line: filter response, dashed line: frequency response of the FIR part, dotted lines: the individual responses of the second-order sections.

6.2. Fixed-pole design example

Here a fixed-pole parallel filter is designed to model a piano soundboard response with 50 second-order sections. The target impulse response is highly nonminimum-phase, as can be seen in Fig. 7 (a), thin line (notice that the peak of the response is at around $n = 200$ and not at $n = 0$ as for minimum-phase systems). Since the second-order sections cannot efficiently model the rising part of the response, a 200th order FIR parallel part is added, as suggested in [6]. As can be seen in the Fig. 6 thick line and Fig. 7 (a) thick line, the filter follows the specification quite well. However, due to the overlap of the FIR and IIR parts, now 120 dB difference arises between the responses of some second-order sections and the final transfer function. This leads to losing 20 bits precision, which makes the implementation of this filter in single precision floating point arithmetic problematic.[1]

The dynamic range problem can be avoided if a parallel filter with a delayed IIR part is used. In the next example,

[1] Note that if instead of the fixed-pole design we estimated a high-order IIR filter to the same piano soundboard response and computed the parallel form via partial fraction expansion, a similarly severe dynamic range problem would arise.

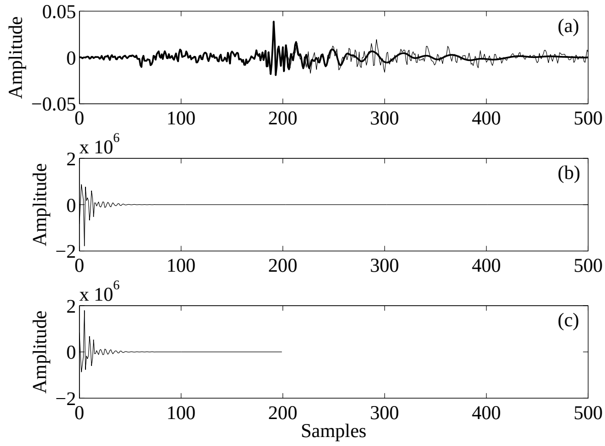

Fig. 7: Time-domain responses of the parallel filter used in Fig. 6: (a) target impulse response (thin line) and filter impulse response (thick line), (b) the response of the second-order sections in a different scale, and (c) the response of the 200th order FIR part.

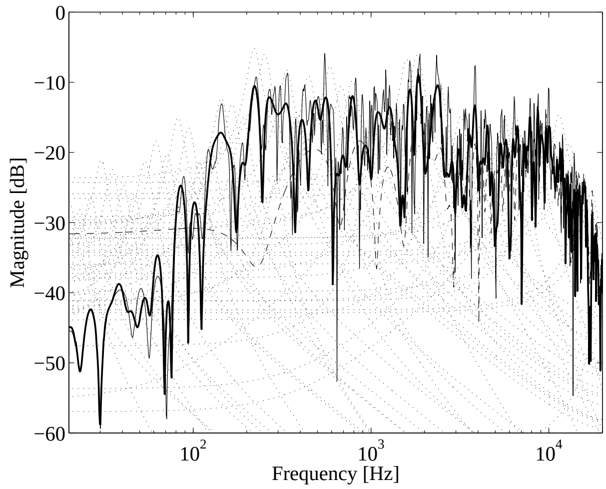

Fig. 8: Fixed-pole parallel filter with a delayed IIR part obtained by the modified design of Sec. 5.3 from the same target as in Fig. 6. Thin line: target frequency response, thick line: filter response, dashed line: frequency response of the FIR part, dotted lines: the responses of the second-order sections.

the filter is designed by the modified method of Sec. 5.3. It can be seen in Fig. 8 that now the gains of the individual sections (dotted lines) and the FIR part (dashed line) are in the same range as that of the total transfer function.

The same responses are displayed in the time-domain in Fig. 9, showing how the delayed IIR (b) and FIR (c) parts are combined to form the total impulse response (a).

7. CONCLUSION

Parallel second-order filters are often used instead of high-order direct form IIR filters because of their improved numerical properties. When the numerator of the original transfer function has equal or higher order than that of the denominator, an FIR part arises in parallel with the IIR sections. Unfortunately, if the parallel form is obtained via the usual partial fraction expansion, the gains of the individual sections can significantly exceed the gain of the total transfer function, leading to the reduction of useful dynamic range. This paper has proposed the use of a modified parallel filter structure in which the IIR sections are delayed so that there is no overlap with the FIR part. The parameters of the delayed parallel filter are either obtained by a simple conversion from the original parallel filter, or from the direct-form IIR filter by an alternate partial fraction expansion, which applies long division on the reversed polynomial when extracting the FIR part. In addition, a modified fixed-

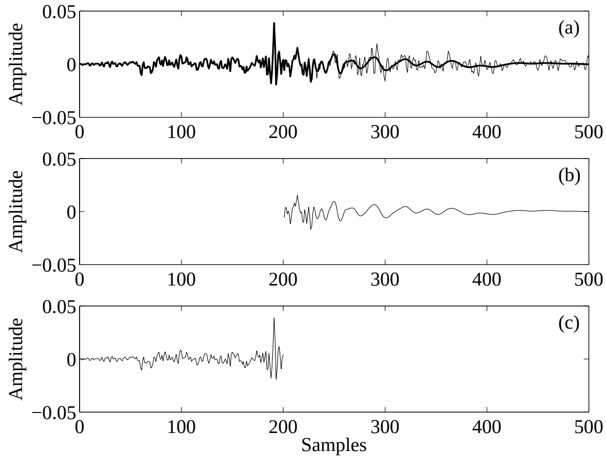

Fig. 9: Time-domain responses of the parallel filter used in Fig. 8: (a) target impulse response (thin line) and filter response (thick line), (b) the response of the second-order sections, and (c) the response of the 200th order FIR part.

pole design procedure has been presented that results in a delayed parallel filter. Besides the numerical benefits, this modified design algorithm also reduces design complexity.

Matlab code for computing a delayed parallel filter from direct-form IIR filters and for the modified fixed-pole design can be found at `http://www.mit.bme.hu/~bank/delparf`

8. ACKNOWLEDGEMENT

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

9. REFERENCES

[1] T. W. Parks and C. S. Burrus, *Digital Filter Design*. USA: John Wiley and Sons, 1987.

[2] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing*. Upper Saddle River, New Jersey, USA: Prentice-Hall, 1999.

[3] W. Chen, "Performance of cascade and parallel IIR filters," *J. Audio Eng. Soc.*, vol. 44, no. 3, pp. 148–158, 1996.

[4] J. O. Smith III, *Introduction to Digital Filters with Audio Applications*. http://ccrma.stanford.edu/˜jos/filters/, Sep. 2007, online book.

[5] A. Krukowski, I. Kale, and G. D. Cain, "Decomposition of IIR transfer functions into parallel arbitrary-order IIR subfilters," in *Proc. IEEE Nordic Signal Processing Symposium (NORSIG'96)*, Espoo, Finland, Sep. 1996.

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

[7] A. Härmä, M. Karjalainen, L. Savioja, V. Välimäki, U. K. Laine, and J. Huopaniemi, "Frequency-warped signal processing for audio applications," *J. Audio Eng. Soc.*, vol. 48, no. 11, pp. 1011–1031, Nov. 2000.

[8] M. Karjalainen and T. 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*, vol. 2007, p. 13, 2007, article ID 60949, doi:10.1155/2007/60949.

[9] B. Bank, "Audio equalization with fixed-pole parallel filters: An efficient alternative to complex smoothing," *J. Audio Eng. Soc.*, vol. 61, no. 1/2, pp. 39–49, Jan. 2013.

[10] P. D. Hatziantoniou and J. N. Mourjopoulos, "Generalized fractional-octave smoothing for audio and acoustic responses," *J. Audio Eng. Soc.*, vol. 48, no. 4, pp. 259–280, Apr. 2000.

[11] B. Bank, "Loudspeaker and room equalization using parallel filters: Comparison of pole positioning strategies," in *Proc. 51st AES Conf. on Loudspeakers and Headphones*, Helsinki, Finland, Aug. 2013.

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

[13] ——, "Logarithmic frequency scale parallel filter design with complex and magnitude-only specifications," *IEEE Signal Process. Lett.*, vol. 18, no. 2, pp. 138–141, Feb. 2011.

[14] K. Steiglitz and L. E. McBride, "A technique for the indentification of linear systems," *IEEE Trans. Autom. Control*, vol. AC-10, pp. 461–464, Oct. 1965.