



Audio Engineering Society

Convention e-Brief 337

Presented at the 142nd Convention
2017 May 20–23, Berlin, Germany

This Engineering Brief was selected on the basis of a submitted synopsis. The author is solely responsible for its presentation, and the AES takes no responsibility for the contents. All rights reserved. Reproduction of this paper, or any portion thereof, is not permitted without direct permission from the Audio Engineering Society.

Quantization noise of warped and parallel filters using single-precision floating-point arithmetic

Balázs Bank¹ and Kristóf Horváth²

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

²*Prolan Process Control Co., Budakalász, Hungary*

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

ABSTRACT

For audio filter and equalizer design it is desirable to take into account the frequency resolution of hearing. Therefore, various specialized filter design methodologies have been developed, from which warped and parallel filters are particularly appealing due to their simple design and good approximation properties. This paper compares the roundoff noise of two different warped IIR implementations (original all-pass and dewarped cascade) with that of fixed-pole parallel filters in single-precision floating-point arithmetic. It is shown by simulations that the parallel filter provides the best compromise between quantization noise and computational complexity, since it significantly outperforms the dewarped series second-order IIR implementation in terms of noise performance, while requires less computational resources compared to the allpass-based warped IIR structure.

1. INTRODUCTION

Designing a filter that models an acoustic transfer function or an equalizer that improves the response of the system is a common task in audio signal processing. Therefore, various audio filter design methodologies have been developed that take into account the frequency resolution of the human auditory system. This paper compares the quantization noise performance of warped filters and fixed-pole parallel filters when using floating-point arithmetic.

2. WARPED FILTERS

The most often used audio filter design technique is frequency warping where the unit delays of traditional FIR or IIR filters are replaced by first-order all-pass filters

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

resulting in the transformation of the frequency axis [1]. This transforms the inherently linear frequency resolution of FIR and IIR filters, and by the careful choice of the warping parameter it is possible to implement a

logarithmic-like resolution similar to the human hearing. The filter design is done using traditional filter design methods (e.g., windowed FIR, Least Squares, Prony, Steiglitz-McBride) based on a transformed (warped) version of the target impulse response. Frequency warping, due to the better allocation of frequency resolution, allows the reduction of the filter order for the same perceptual quality in comparison to traditional FIR and IIR filters. Unfortunately, FIR and IIR filters require specialized allpass-based filter structures [1]; therefore the reduction in filter order does not translate to the same reduction in terms of required arithmetic operations. This is because the special warped IIR structure requires around two times larger computational complexity compared to direct-form IIR filters having the same order [1]. This can be overcome by converting the warped FIR or IIR filters to a cascade of second-order sections [2].

3. PARALLEL FILTERS

Kautz filters can be seen as a generalization of warped FIR filters where the warping parameter can be different for all the sections [3]. Kautz filters model the target impulse response as a linear combination of orthonormal basis functions, leading to a simple design procedure based on a scalar product. Their only drawback is that a special series-parallel structure is needed for their implementation, leading to larger complexity compared to direct-form IIR filters of the same order.

By giving up the orthonormality of Kautz filters, the idea of fixed-pole parallel filters was born [4]. Parallel filters aim to approximate the target impulse response as a linear combination of exponentially decaying sinusoidal functions, implemented as a parallel set of second-order IIR filters and an optional FIR path:

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

where K is the number of second order sections. It has been shown in [5] that the resulting frequency response is practically equivalent to Kautz filters given the same pole locations, while the number of required arithmetic operations is reduced to two-third.

The steps of designing parallel filters are the following: First the pole locations are determined. This can be either done by manually fixing the poles directly influencing

the frequency resolution of the design (e.g., setting them to a logarithmic scale results in a logarithmic frequency resolution), or by taking the poles from a supplementary IIR design. This usually means designing a warped IIR filter based on the target response. (For different pole positioning methods, see [6, 7]). The poles determine the denominators of the parallel second-order sections, and now Eq. (2) becomes linear in the numerator coefficients that can be estimated in a closed form by the least squares equations [5].

If the poles of the parallel filter are obtained based on a warped IIR design, the frequency response of the parallel filter will be practically the same as that of the warped IIR filter [4]. Therefore, the question arises, is there any particular advantage of the parallel filter implementation. In this engineering brief the quantization performance of these filter implementations are compared when using single-precision floating-point arithmetic.

4. COMPARISON METHODOLOGY

While the analytical derivation of quantization noise level is possible for floating-point arithmetic (see, e.g., [8]), it is significantly more complicated compared to the fixed-point case due to the fact that the amplitudes of noise sources are signal dependent. Therefore, the comparison is done by using MATLAB simulations where the quantization noise is computed as the difference of the tested (single-precision floating-point) and reference (double-precision floating-point) filter implementations. The coefficients of the reference implementations are truncated to single-precision so that the two versions have the same transfer function.

As for the excitation signal, pink noise is used due to its spectral shape more similar to audio signals in comparison to the commonly used white noise excitation. The quantization noise is analyzed in third-octave bands common in audio signal processing.

Three filter implementations will be compared:

- Warped IIR filter in the all-pass form of Fig. 6 in [1], typically requiring two times more computational resources compared to the other two options
- Warped IIR filter converted to series second-order sections [2] implemented using DF1 structures
- Fixed-pole parallel filter with a design based on the warped poles [4], implemented by DF1 structures




Fig. 1: Loudspeaker-room response equalization: (a) the smoothed loudspeaker-room response, and (b) equalized by a 40th order filter. The target specification is displayed by dashed lines. The transfer function of the equalizer is displayed by curve (c) and the pole frequencies are displayed with crosses. The curves are offset for clarity.

5. DESIGN EXAMPLES

5.1. Room equalization

First, a loudspeaker-room equalization example is presented, taken from [5]. Here the smoothed system response is used, and a 40th order WIIR equalizer is designed using $\lambda = 0.95$, corresponding to the example of Fig. 3 (d) in [5]. The smoothed loudspeaker-room response is shown in Fig. 1 (a), while the equalized transfer function in (b), and finally, the frequency response of the equalizer filter is displayed in Fig. 1 (c). From the WIIR filter a series second-order implementation is obtained by dewarping [2], and a parallel filter is designed based on the poles of the warped IIR filter [4]. These two filters have the same transfer function as the warped IIR filter they are derived from. However, since their implementation is different, we expect different quantization noise performance.

The noise levels in third-octave bands for the three implementations are displayed in Fig. 2. The third-octave levels of the useful output signal are displayed by dash-dotted line for reference: the SNR in the various bands

is simply the difference of this dashed line and the noise level curves. We note here that on the contrary to fixed-point quantization, the SNR value is independent of the absolute level of the input signal since the quantization noise is level dependent: for two times larger input the noise increases by two times as well, leaving the SNR intact. This is a very useful property for audio applications since for low signal levels the noise will be decreased by a similar amount and thus it is more likely to be masked by the program material.

The dotted line in Fig. 2 is the quantization noise level when only the output is quantized to single-precision: this can be considered as the best case that can be theoretically achieved.

It can be observed in Fig. 2 that the all-pass warped IIR implementation (thin solid line) has the lowest noise at low frequencies, while its performance gets worse in the mid and high frequency range. The opposite can be observed for the dewarped cascade (dashed line) and for the parallel filter (thick solid line) implementations: actually this behavior is more preferable since the audibility threshold is higher at low frequencies, making this kind of noise profile less disturbing. It can also be seen that the parallel filter has the lowest quantization noise levels in general.

5.2. Soundboard modeling

Next, a piano soundboard modeling example is presented (see [4] for details). The force-pressure frequency response of the measured piano soundboard is shown in Fig. 3 (a), which is modeled by a 200th order warped IIR filter ($\lambda = 0.8$) displayed in (b). Again, from the WIIR filter a series second-order implementation is computed by dewarping [2], and a parallel filter is obtained by using the poles of the warped IIR filter [4], giving practically identical transfer functions.

The noise levels in third-octave bands for the three implementations are shown in Fig. 4. It can be seen in Fig. 2 that the warped IIR implementation (thin solid line) has again the lowest noise at low frequencies, while its performance gets worse in the mid and high frequency range: they cross with the parallel implementation (thick solid line) at 1 kHz. Similarly to the previous example, the dewarped series implementation has the worst performance (dashed line).

The signal-to-noise ratios for the two example designs are shown in Table 1, obtained by summing the signal

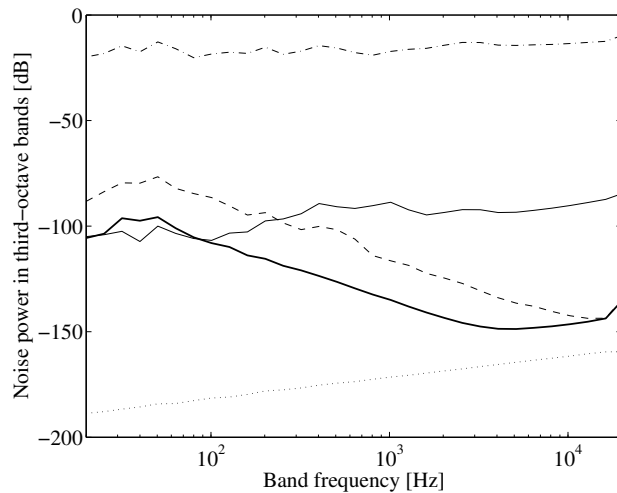


Fig. 2: Quantization noise levels in third-octave bands for the transfer function of Fig. 1 (c) implemented in single-precision floating-point arithmetic as a response to pink-noise excitation with unity power. The thin solid line is for the WIIR all-pass implementation, the dashed line is for the WIIR filter dewarped to cascade form, and the thick solid line is for the parallel filter. For reference, the dash-dotted line on the top shows the output signal levels in the bands, while the dotted line in the bottom corresponds to performing the filtering in double-precision and quantizing only at the output.

and noise powers in the third-octave bands. It can be noticed that the parallel filter implementation shows a consistent 20 dB improvement over the dewarped series version, while it has the same computational complexity. For the soundboard modeling case, the parallel filter is outperformed by the WIIR all-pass implementation, at the price of roughly doubled computational complexity. Interestingly, for the room EQ case the WIIR implementation leads to significantly larger noise levels, the reasons for this difference are yet to be understood.

6. CONCLUSION

This paper has compared the quantization noise performance of three filter implementations based on a warped IIR design using single-precision floating-point arithmetic: the allpass-based WIIR structure of [1], dewarped to series second-order sections using the formulas of [2] and the fixed-pole parallel filter using the poles of the WIIR filter [4]. These three implementations have the




Fig. 3: Piano soundboard modeling: (a) the measured force-pressure response, and (b) modeled by a 200th order filter. The pole frequencies are displayed by crosses. The curves are offset for clarity.

	WIIR	Dewarped	Parallel
Room EQ	77.3 dB	71.8 dB	90.1 dB
Piano	116.1 dB	62.4 dB	83.0 dB

Table 1: Signal-to-noise ratios computed for the room equalization example of Fig. 2 and for the piano soundboard modeling example of Fig. 4.

same frequency response but their quantization noise behavior differ. The examples show that the dewarped series implementation has the worst performance and parallel filters show a 20 dB improvement in comparison. The WIIR all-pass implementation gives a varying result, understanding the reason for this is left for future research.

To sum up, the parallel filter has similarly low computational complexity as the dewarped series implementation, but it provides better noise performance. The warped IIR all-pass implementation improves the noise performance for one of the design cases, at the price of a more complicated filter structure. Therefore the parallel filter seems to be the best compromise among the three options discussed here.

Additional advantages of parallel filters over warped IIR




Fig. 4: Quantization noise levels in third-octave bands for the transfer function of Fig. 3 (b) implemented in single-precision floating-point arithmetic as a response to pink-noise excitation with unity power. The thin solid line is for the WIIR all-pass implementation, the dashed line is for the WIIR filter dewarped to cascade form, and the thick solid line is for the parallel filter. For reference, the dash-dotted line on the top shows the output signal levels in the bands, while the dotted line in the bottom corresponds to performing the filtering in double-precision and quantizing only at the output.

filters (either in the original form or dewarped) include the flexibility of choosing the frequency resolution by using alternative design methods (e.g., using a predetermined pole set [5]), and the possibility of full code parallelization yielding significant performance benefits in modern architectures [9].

Future research includes a more systematic testing of the three filter implementations with a larger variety of design examples and understanding the underlying reasons for the varying performance, especially that of the allpass-based WIIR filter. Furthermore, incorporating the ITU-R 468 weighting in the comparison would better reflect the noise sensitivity of human hearing.

7. ACKNOWLEDGEMENT

The work of B. Bank has been supported by the ÚNKP-16-4-III New National Excellence Program of the Hungarian Ministry of Human Capacities.

8. REFERENCES

- [1] M. Karjalainen, A. Härmä, and U. K. Laine, “Realizable warped IIR filters and their properties,” in *Proc. IEEE Int. Conf. Acoust. Speech and Signal Process.*, vol. 3, Munich, Germany, Apr. 1997, pp. 2205–2208.
- [2] S. S. Ahuja and S. C. D. Roy, “Variable digital filters,” *IEEE Trans. Circ. Syst.*, vol. CAS-27, no. 9, pp. 836–838, Sep. 1980.
- [3] T. Paatero and M. Karjalainen, “Kautz filters and generalized frequency resolution: Theory and audio applications,” *J. Audio Eng. Soc.*, vol. 51, no. 1–2, pp. 27–44, Jan./Feb. 2003.
- [4] 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>.
- [5] —, “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.
- [6] —, “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.
- [7] E. Maestre, G. P. Scavone, and J. O. Smith, “Design of recursive digital filters in parallel form by linearly constrained pole optimization,” *IEEE Signal Process. Lett.*, vol. 23, no. 11, pp. 1547–1550, Nov. 2016.
- [8] E. P. F. Kan and J. K. Aggarwal, “Error-analysis of digital filter employing floating point arithmetic,” *IEEE Trans. Circ. Theory*, vol. CT-18, no. 6, pp. 678–686, Nov. 1971.
- [9] J. A. Belloch, B. Bank, L. Savioja, A. Gonzalez, and V. Välimäki, “Multi-channel IIR filtering of audio signals using a GPU,” in *Proc. IEEE Int. Conf. Acoust. Speech and Signal Process.*, Florence, Italy, May 2014, pp. 6692–6696.