Audio Engineering Society

Convention e-Brief

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



# Warped implementations of parallel second-order filters with optimized quantization noise performance

Balázs Bank[1] and Kristóf Horváth[2]

[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

Fixed-pole second-order parallel filters provide an efficient way of implementing IIR filters with a logarithmic frequency resolution. However, the fine frequency resolution needed at low frequencies can only be achieved by poles near the unit circle. This may lead to large roundoff noise at low frequencies when the filters are implemented using bit-depths of 24 bits or lower in fixed-point arithmetic. This paper investigates the performance improvement when the parallel second-order sections are implemented as warped IIR filters. In addition, an analytical expression is given for computing the warping parameter as a function of the pole location of the original second-order section so that the quantization noise power is minimized.

## 1. INTRODUCTION

Fixed-pole parallel filters provide an efficient way of implementing IIR filters with a flexible allocation of frequency resolution [1]. The basic idea of parallel filters is that the transfer function is composed of a parallel set of second-order IIR filters having two poles and one zero, and an optional FIR path, given as

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

where $K$ is the number of second order sections.

Traditionally, parallel second-order filters are obtained from high-order IIR filters by partial fraction expansion [2]. However, in the methodology of designing parallel filters directly, it is important that the poles are predetermined. Fixing the poles leads to a linear-in-parameter problem since only the numerator coefficients need to be optimized in Eq. (1), which can be obtained by a least squares fit in closed form [1]. An advantage of fixing the poles is that this way we obtain a direct control over the frequency resolution of the filter design; setting the pole frequencies according to a logarithmic scale results in a logarithmic frequency resolution, but of course applying

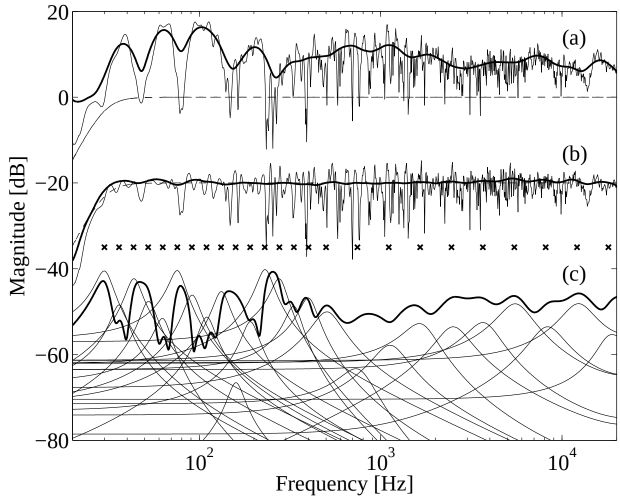


**Fig. 1:** Loudspeaker-room response equalization: (a) unequalized loudspeaker–room response, and (b) equalized by a parallel filter having 25 second-order IIR sections. The thick lines show the third-octave smoothed versions of the transfer functions, and the target specification is displayed by dashed lines. The transfer function of the equalizer is displayed in (c) by a thick line, while the magnitude responses of the second-order sections are shown by thin lines. The pole frequencies are displayed with crosses. The curves are offset for clarity.

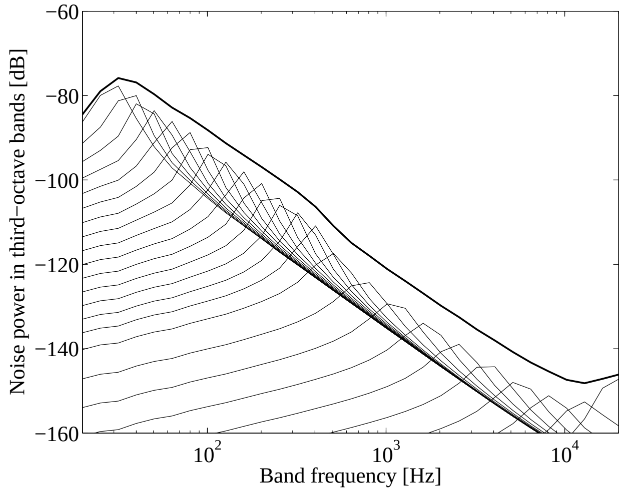


**Fig. 2:** Quantization noise levels in third-octave bands for the transfer functions of Fig. 1 (c) implemented using DF1 structure in 24 bit fixed-point arithmetic. Thin lines show the noise levels $P_k(f_n)$ of the second-order sections, while the total noise power $P(f_n)$ in the third-octave bands is displayed by a thick line as a function of band frequency $f_n$.

different resolution in various regions of the frequency range is also possible.

For example, in loudspeaker-room equalization we may wish to equalize the low-frequency room modes at a finer detail compared to the high-frequency response of the system. This is displayed in Fig. 1 where we have higher pole density in the problematic region of room modes (below 500 Hz) compared to mid and high frequencies, as shown by the crosses indicating the pole frequencies of the second-order sections. It can be seen in Fig. 1 (c) that the parallel equalizer provides a smooth overall response, without attempting to counteract the sharp notches of the transfer function.

Note that besides using a predetermined pole set it is also possible to allocate the poles by automatic procedures, see [3] for the comparison of available methods.

## 2. QUANTIZATION NOISE PERFORMANCE

Independently of the pole positioning method, the fine frequency resolution needed at low frequencies for a logarithmic scale can only be achieved by poles near the unit circle. If the second-order sections are implemented as traditional direct form filters, this means that their quantization noise will be boosted by either the all-pole part or the complete transfer function, depending on the type of implementation [2].

Figure 2 shows the quantization noise levels in third-octave bands for the above equalizer as a function of the center frequency of the third-octave bands for Direct Form 1 (DF1) implementation using 24 bit fixed-point (fractional) arithmetic.

The noise levels are computed analytically by assuming a typical DSP architecture (quantization at the accumulator), and modeling the quantization effects as uncorrelated additive white noise having the standard deviation $\sigma_n = 2^{-b+1}/\sqrt{12}$, where $b$ is the number of bits [2]. In DF1 implementation there is only one summation (and thus quantization) point, and the output noise power spectral density $E_k(\vartheta)$ of the sections is simply the quantization noise power spectral density multiplied by the square of the all-pole transfer functions, as shown in Eq. (2):

$$E_k(f) = \left| \frac{1}{1 + a_{k,1}e^{-j\vartheta} + a_{k,2}z^{-j2\vartheta}} \right|^2 \frac{\sigma_n^2}{f_\mathrm{s}/2}, \qquad (2)$$

where $\vartheta$ is the angular frequency $\vartheta = 2\pi f/f_\mathrm{s}$, with $f_\mathrm{s}$ being the sampling rate.

Power spectral densities are not easy to interpret; on the contrary, third-octave noise analysis is common in assessing the performance of audio systems. The noise power in third-octave bands is obtained by (numerically) integrating the power spectral densities as

$$P_k(f_n) = \int_{f_n/c}^{f_n c} E_k(f)df, \qquad (3)$$

where $c = 2^{1/6}$ corresponds to a sixth-octave distance from the band center $f_n$. The total noise power $P(f_n)$ in the third-octave band centered at $f_n$ is simply the sum of $P_k(f_n)$ for all $K$, displayed by thick line in Fig. 2.

As a reference, a full-scale sine wave has the power of $-3$ dB, and the total noise power summed for all the third-octave bands is $-70.7$ dB, leading to an $SNR = 67.7$ dB. This is on the edge of being audible, since a large part of the noise power is coming from low frequencies, where the audibility threshold is higher. Of course it is practical to have some headroom so typical program material will have less power than $-3$ dB, decreasing the signal-to-noise ratio. Also, with lower bit-depths, the corresponding curves of Fig. 2 are shifted: a 20 bit implementation would mean all noise levels moved up by $24$ dB.

On the other hand, when using 32 bit arithmetic, the noise levels are moved by $-48$ dB, so it is unlikely that quantization noise will cause any performance degradation.

## 3. WARPED IMPLEMENTATION

A common solution to fight against quantization noise is to implement the second-order sections by special filter structures (e.g., Kingsbury or Zölzer) instead of the usual direct or transposed forms [4]. Of course this leads to an increase of computational complexity; therefore it is suggested to implement only the problematic low frequency sections in a special form.

This paper investigates the performance improvement when the problematic second-order sections are implemented as warped IIR filters. In warped IIR filters, the unit delays of traditional IIR filters are replaced by the all-pass filter

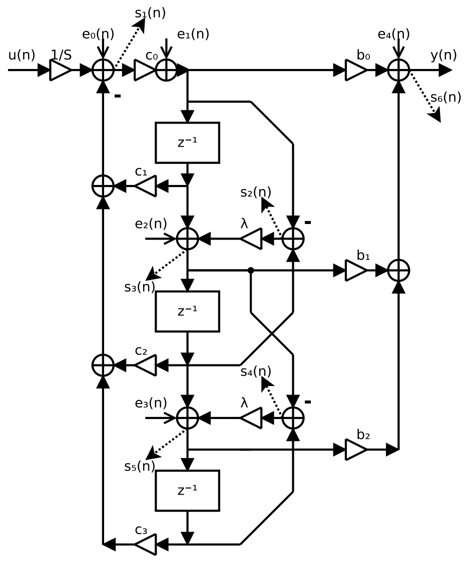


**Fig. 3:** Second-order warped IIR structure based on [5]. The $e_i(n)$ are the independent noise sources modeling quantization noise, and $s_i(n)$ are the impulse responses that are used for computing the necessary scaling $1/S$.

$$D(z) = \frac{z^{-1} - \lambda}{1 - \lambda z^{-1}} \qquad (4)$$

where the warping parameter $\lambda$ allows the distortion of the frequency axis [5]. Because of these additional all-pass filters, warped IIR implementations require special structures and thus around two times more computational resources compared to direct from IIR filters [5]. Therefore, warped IIR filters are often converted to direct-form filters [5], or to series or parallel second-order sections [6]. Here we make the opposite: we convert the direct form second-order sections to warped implementations using the equations of [6] with a negative warping parameter $-\lambda$. Now the question arises how this influences the quantization noise performance.

Figure 4 shows the third-octave-band noise levels when the second-order sections are implemented by warped IIR structures displayed in Fig. 3 with $\lambda = 0.9$. The noise levels were computed similarly to the DF1 case with the difference that now an input scaling $1/S$ is applied so that none of the summation points $s_i(n)$ of the warped IIR structure go to overload. This is because in

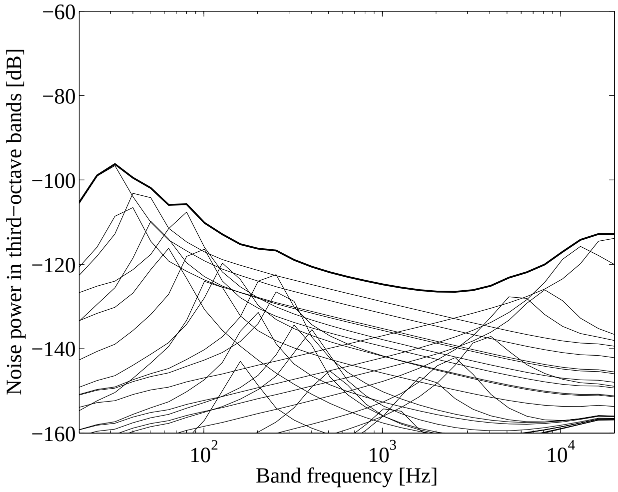


**Fig. 4:** Quantization noise levels in third-octave bands for the transfer functions of Fig. 1 (c) implemented using second-order warped IIR sections in 24 bit fixed-point arithmetic with $\lambda = 0.9$. Thin lines show the noise levels $P_k(f_n)$ of the second-order sections, while the total noise power $P(f_n)$ in the third-octave bands is displayed by a thick line as a function of band frequency $f_n$.

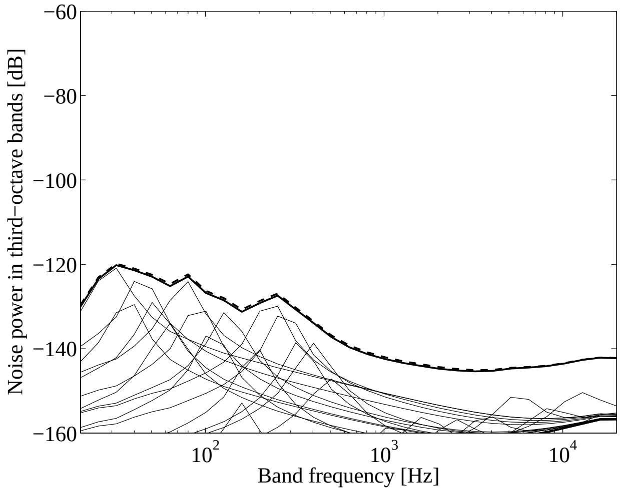


**Fig. 5:** Quantization noise levels in third-octave bands for the transfer functions of Fig. 1 (c) implemented by second-order warped IIR sections in 24 bit fixed-point arithmetic using optimized $\lambda_k$ for all sections. Thin lines show the noise levels $P_k(f_n)$ of the second-order sections, while the total noise power $P(f_n)$ in the bands is displayed by a thick line as a function of band frequency $f_n$. The dashed line shows the total noise power when the analytical $\tilde{\lambda}_k$ parameters obtained from Eq. (5) are used instead of the numerically optimized $\lambda_k$ values.

WIIR structures internal overflow can happen even if the total transfer function of the section is below 0 dB at all frequencies. In addition, since there are more quantization points, the output noise powers of the sections are computed as a sum of independent noise sources $e_i(n)$ all filtered by the square of the respective transfer functions. The noise transfer functions are computed numerically by taking the Fourier-transform of the impulse responses from the quantization noise source $e_i(n)$ to the output.

It can be seen in Fig. 4 that for $\lambda = 0.9$ noise performance improves radically at low frequencies, while worsens in the high frequency range. The total noise power is $-91.7$ dB, leading to a $20$ dB improvement compared to the DF1 case.

## 4. OPTIMIZING NOISE PERFORMANCE

There is no practical constraint forcing the $\lambda$ values being the same for all sections: next we investigate the effect of having individual $\lambda_k$ values. For this, an optimization routine has been developed that searches for the $\lambda_k$ parameter for each second-order section so that the total noise power (sum of all third-octave bands) is minimized for that section. The obtained noise performance for the same room equalizer is displayed in Fig. 5.

It can be observed that the noise levels are decreased radically due to the optimization, the total noise power being $-113.5$ dB. This means that even implementing the equalizer using 20 bit arithmetic (curves shifted up by 24 dB) would give acceptable noise levels.

The $\lambda$ values found by the optimizer are displayed in Fig. 6 solid line. The question arises whether there is any underlying reason that explains why a particular lambda value produces the lowest noise for a specific pole frequency. By observing the poles after frequency warping, it turns out that the warped pole angle is almost 90 degrees for all cases. It can be shown that the warped pole $\tilde{p}_k$ has exactly 90 degrees angle (zero real part) when the following quadratic equation is satisfied:

$$\tilde{\lambda}_k^2 \mathrm{Re}\{p_k\} + \tilde{\lambda}_k(-1 - |p_k|^2) + \mathrm{Re}\{p_k\} = 0. \quad (5)$$

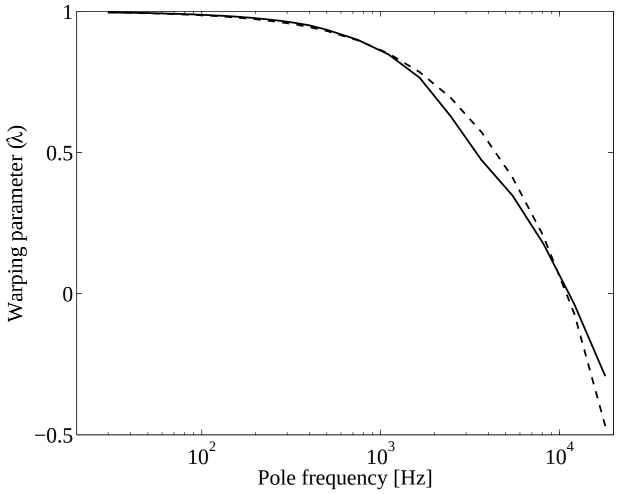


**Fig. 6:** Optimized $\lambda_k$ values used for computing the noise powers of Fig. 5 as a function of the analog pole frequencies of the second-order sections. Solid line: numerically optimized values, dashed line: analytic $\tilde{\lambda}_k$ parameters computed using Eq. (5).

From the resulting two roots, the one with $|\tilde{\lambda}_k| \leq 1$ has to be chosen. Figure 6 dashed line displays the analytical $\tilde{\lambda}_k$ values computed by solving Eq. (5), showing a good match.

Even more convincing is the total noise level computed using the analytical $\tilde{\lambda}_k$ values, shown by dashed line in Fig. 6, being almost indistinguishable from the numerically optimized noise performance (solid line). Similarly accurate match has been observed for other design examples, justifying the use of Eq. (5) for computing the optimal warping parameter of second-order sections instead of a more complex noise optimization routine.

## 5. CONCULSION

This paper has shown that a significant improvement in quantization noise performance of second-order IIR filters can be achieved when implementing them as warped IIR structures. For the example showed in this paper, more than $40$ dB improvement has been achieved, and similar improvements have been observed for other design cases. This comes at a price of larger computational complexity, however, this might be outweighed by the fact that the need for larger bit-depth is eliminated. In addition, it is suggested that only the problematic sections are implemented as warped IIR filters. In Fig. 2 the crossover frequency would be at around 200-300 Hz.

In addition to showing the roundoff noise benefits using warped implementations, an analytical formula for computing the optimal lambda parameter as a function of the original pole location has been given, showing negligible performance loss compared to the numerically optimized value.

We note that warped implementations can improve the noise performance of parallel and series IIR filters in the same way when they are computed by factoring high-order transfer functions, not only for the fixed-pole design used as an example in this paper.

Future research includes performing the optimization by using ITU-R 468 and A noise weighting that would better reflect the audibility of quantization noise.

## 6. 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.

## 7. REFERENCES

[1] 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.

[2] A. V. Oppenheim, R. W. Schafer, and J. R. Bruck, *Discrete-Time Signal Processing*. Englewood Cliffs, New Jersey, USA: Prentice-Hall, 1975.

[3] 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.

[4] U. Zölzer, *Digital Audio Signal Processing*. West Sussex, England: John Wiley and Sons, 1997.

[5] 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.

[6] 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.