TRUNCATION NOISE ANALYSIS OF NOISE SHAPING DSP SYSTEMS WITH APPLICATION TO CIC DECIMATORS

Vesa Lehtinen and Markku Renfors

Institute of Communications Engineering, Tampere University of Technology
P.O.Box 553, FIN-33101, Tampere, Finland
Vesa.Lehtinen@tut.fi, Markku.Renfors@tut.fi

***Abstract* – In ordinary digital filters, reducing the wordlength of a signal in truncation always increases the quantization error variance at system output. We show that this law is not always valid if quantization noise shaping, e.g., error feedback, is used. This fact should be taken into account in the design of wordlength optimization algorithms.**

Furthermore, we show that cascaded integrator-comb (CIC) decimators provide inherent noise shaping. This is a consequence of pole-zero cancellation at passband, which has become possible through the use of modulo arithmetic. Thus, the wordlengths in the comb stages can be kept small, which can also reduce the complexity of the successive filter stage(s) usually following a CIC decimator.

***Keywords* – CIC decimator, error feedback, noise shaping, quantization noise, truncation, wordlength optimization**

1 INTRODUCTION

In [1], a truncation noise model for DSP systems was derived, giving more reliable estimates than earlier models. This more accurate model reveals a previously undiscovered property of truncation noise, having an effect on signal wordlength optimization.

The structure of this paper is the following. In Section 2, the CIC decimator structure is briefly reviewed. In Section 3, the truncation noise model is given, and a formula of the contribution of multiple truncations in a cascade of digital filters to the overall output noise is derived. A new phenomenon is discovered, having importance especially in optimized ASIC and FPGA implementations where wordlengths can be chosen freely and truncations of just a few bits can be made to save resources. In Section 4, we show that CIC filters provide inherent noise shaping for truncation errors generated within the CIC structure. CIC filters are also used as an example of the theory derived in Section 3 and in the simulations used to prove the theory.

2 CIC DECIMATORS

The cascaded integrator-comb (CIC) decimator [2] is an efficient recursive implementation form of first- or higher-order running sum filters. The transfer function and frequency response of a CIC decimation filter of N^{th} order are

$$H(z) = \left(\frac{1-z^{-R}}{1-z^{-1}}\right)^N \tag{1}$$

and

$$H(e^{j\omega}) = e^{-jN\frac{R-1}{2}\omega} \lim_{\theta \to \omega} \left(\frac{\sin(R\theta/2)}{\sin(\theta/2)}\right)^N, \tag{2}$$

respectively, where R is the decimation factor. In Figure 1, the block diagram of a second-order CIC decimator is presented. By exploiting modulo arithmetic, the denominator part (integrators) of the transfer function can be implemented separately from the numerator part (combs). This reduces the addition rate of the numerator part (combs) and the amount of memory required. Integrators are unstable recursive filters, but because of modulo arithmetic, the overflows occurring in the integrators are cancelled by the combs, provided that all substages have equal level MSBs (i.e., have equal moduli) and the wordlength of the first integrator is sufficient.

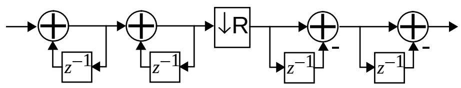

Figure 1. Block diagram of a second-order CIC filter.

3 TRUNCATION NOISE MODEL

Truncation is a nonlinear process but can be modeled as adding independent, evenly distributed white noise into the signal.

Let us define the *LSB level* λ of a signal such that the signal can take values $k2^{\lambda}$, where k is an integer within a range limited by the wordlength and number representation. According to [1], the mean and variance of noise introduced in signal wordlength truncation are

$$\mu_T = \frac{1}{2}(2^{\lambda_{\text{before}}} - 2^{\lambda_{\text{after}}}) \tag{3}$$

and

$$\sigma_T^2 = \frac{1}{12}(4^{\lambda_{\text{after}}} - 4^{\lambda_{\text{before}}}), \tag{4}$$

respectively, where $\lambda_{\text{after}} \le \lambda_{\text{before}}$, and λ_{before} and λ_{after} are the LSB levels before and after the truncation, respectively. In this paper, we concentrate on truncation and ignore rounding because they differ only by the DC offset introduced (average error); truncation requires no hardware and in linear time-invariant (LTI) systems all known DC offset(s) can be removed from the signal with arbitrary precision with a single subtraction of a constant at the system output or (almost) anywhere in the system.

In a cascade of filters $H_s(z)$, $s=1,...,S$, the overall output noise variance is

$$\sigma^2 = \sum_{s=1}^{S+1} \frac{G_s}{12}(4^{\lambda_s} - 4^{\lambda_{s-1} + \lambda_{s-1}^{(c)}}) \tag{5}$$

where λ_s and $\lambda_s^{(c)}$ ($s=1,...,S$) are the (effective) input and coefficient LSB levels of the s^{th} filter stage, respectively, λ_0 and λ_{S+1} are the input and output LSB levels, respectively, $\lambda_0^{(c)} = 0$, and

$$G_s = \begin{cases} \dfrac{1}{2\pi}\displaystyle\int_{-\pi}^{\pi} \left| \prod_{k=s}^{S} H_s(e^{j\omega}) \right|^2 d\omega, & s = 1, ..., S \\ 1, & s = S+1 \end{cases} \tag{6}$$

is the power gain for the s^{th} truncation with G_1 and G_{S+1} corresponding to truncation of the input and output, respectively.

We can rewrite Eq. (5):

$$\sigma^2 = -\frac{G_1}{12}4^{\lambda_0} + \sum_{s=1}^{S} \frac{G_s - 4^{\lambda_s^{(c)}} G_{s+1}}{12} 4^{\lambda_s} + \frac{G_{S+1}}{12} 4^{\lambda_{S+1}} \tag{7}$$

In ordinary filters with no noise shaping, the term

$$\hat{G}_s = G_s - 4^{\lambda_s^{(c)}} G_{s+1} \tag{8}$$

is always positive, because a high attenuation requires small values of $\lambda_s^{(c)}$. Therefore, increasing any λ_s will increase the overall noise variance at system output.

However, if noise shaping, e.g., error feedback in truncation, is used (or occurs implicitly as will be shown in Section 4), it is possible that $\hat{G}_s < 0$ because G_s becomes much smaller with respect to G_{s+1}. In such a case, decreasing λ_s (increasing the wordlength at the input of the s^{th} filter stage) would *increase* the total truncation noise variance at the system output. Therefore, it would be optimal in terms of *both* noise variance and complexity to have $\lambda_s = \lambda_{s+1}$ if other wordlengths are fixed. If $\hat{G}_s = 0$, the wordlength λ_s has no effect on the output noise variance.

The discovery of this phenomenon was made possible by the introduction of the truncation error model based on discrete error distribution [1] because of its negative term in Eq. (4). It was not visible in the traditional model

$$\sigma^2 = \frac{\Delta^2}{12} = \frac{1}{12}4^{\lambda_{\text{after}}}, \tag{9}$$

where Δ is the quantization step. This model was originally derived for analog-to-digital conversion and worked well also in ordinary DSP systems with coefficient wordlengths and truncations larger than few bits, and no truncation noise shaping. Actually, Eq. (9) and Eq. (4) become equivalent if $\lambda_{\text{before}} = -\infty$, i.e., when quantizing an analog signal.

In systems more complex than a cascade of filters, the same principles as explained above can be used. G_s 's are calculated for each pair of truncation noise source and output node, taking all signal paths into account.

In the next section, the theory presented above is applied to and verified with CIC decimators.

4 NOISE SHAPING IN CIC DECIMATORS

The signal wordlength required in a high-order CIC filter may become much larger than that needed at the output. Therefore, wordlength(s) must be reduced. Truncation can be applied inside the CIC structure because it does not affect the overflow cancellation that occurs at the comb stages [2].

Truncation error has a non-zero mean. This is seen as a DC offset in the signal. Because the combs have a zero gain at DC, they cancel the DC offsets introduced by all truncations within the CIC filter.

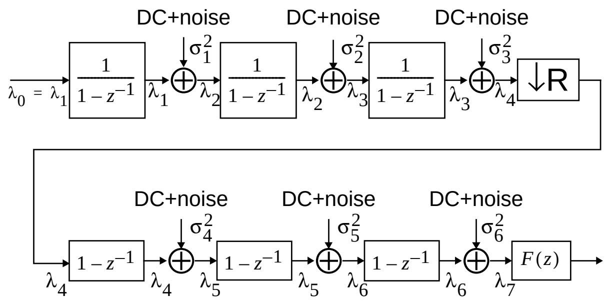

Figure 2. Truncation noise model of a third-order CIC decimator.

Figure 2 depicts the truncation noise model of a CIC filter. It can be seen that truncation noise from all but the last truncation is filtered by a number of (integrator and) comb stages. The CIC decimator is followed by a lowpass filter *F(z)*. Notice that CIC filters are seldom used with no successive filtering because of their passband droop and limited stopband width. The transfer functions of the subfilters can be written as

$$H_s(z) = \begin{cases} \dfrac{1-z^{-R}}{1-z^{-1}}, & s = 0, ..., N \\ 1-z^{-R}, & s = N+1, ..., 2N \\ F(z^R), & s = 2N+1 \end{cases} \quad (10)$$

at the input sample rate.

The system of Figure 2 was simulated to verify the model given in Section 3. A sinusoid with white dithering noise was used as input signal. From Figure 3 it can be seen that the model is quite accurate and predicts well the spectral shape of the filtered noise. In the same simulation it was also verified that it is possible for $\hat{G}_s$ to obtain negative values and, in such a case, it is optimal to use as large a λ_s as possible, i.e., truncate as much as possible before the s^{th} stage. This disproves the conventional assumption that increasing the LSB level would always increase the truncation noise variance at system output.

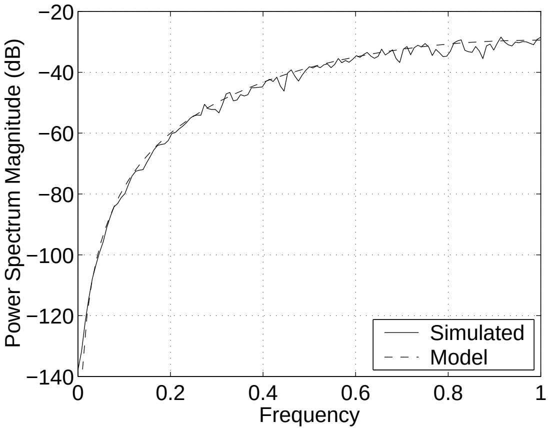

Figure 3. Simulated and theoretical power spectrum of truncation noise generated between the last integrator and the first comb of a third-order CIC decimator and measured at the output of the CIC decimator.

Figure 4 illustrates the effect of the cutoff frequency and stopband attenuation of the lowpass filter $F(z)$ to the value of $\hat{G}_4$ (and thereby, noise shaping) in the system shown in Figure 2. In order to simplify the analysis, a partially ideal lowpass filter *F(z)* was used, having a finite, constant stopband attenuation r_s, an ideal passband with unity gain, and zero-width transition band. The following four facts can be observed: 1) At low cutoff frequencies, changes of the cutoff frequency do not affect the total output noise variance because the cutoff frequency is deep in the noise notch. The noise variance is determined entirely by the stopband attenuation of *F(z)*, and noise shaping gives no benefit. 2) At high cutoff frequencies, shaped noise is let through within the passband, and the benefit of noise shaping is lost. 3) In between, $\hat{G}_4 < 0$, thus a larger λ_4 gives a smaller noise variance at the output of *F(z)*. 4) When $\hat{G}_4 = 0$, λ_4 has no effect on noise variance at the output of *F(z)*.

A typical optimized wordlength profile of a CIC decimator is shown in Figure 5. The first integrator is bound to the wordlength[1] w_1 determined by the DC gain of the CIC filter [2]:

$$w_1 \geq \lceil N \log_2 R \rceil + w_{in}, \quad (11)$$

where w_{in} is the input wordlength.

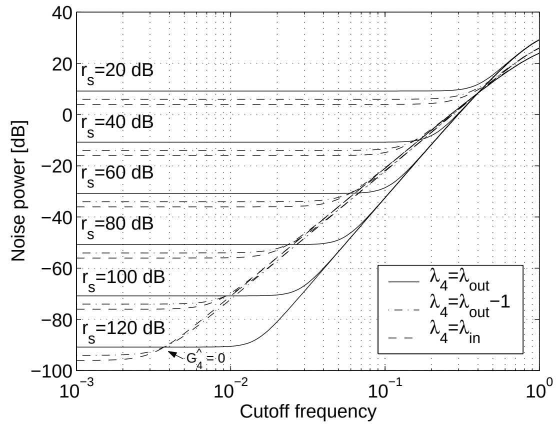

Figure 4. Theoretical output noise variance of the system of Figure 2, plotted as a function of the cutoff frequency of *F(z)* with different stopband attenuations of *F(z)* and wordlengths of the first comb stage. *F(z)* is a partially ideal lowpass filter having a constant stopband attenuation r_s, an ideal passband with unity gain, and zero-width transition band.

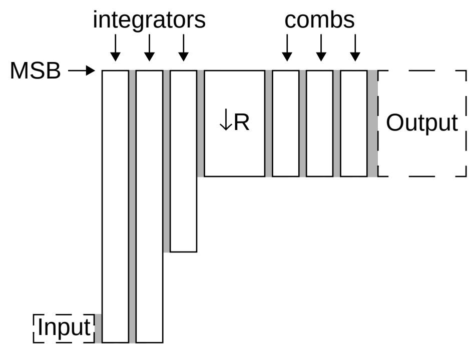

Figure 5. A typical optimized wordlength profile of a CIC decimator. All comb stages have the same wordlength.

Few or no bits at all can be truncated at the next few integrators if the input wordlength is small (e.g., the output of a sigma-delta modulator). Then the wordlength decreases rapidly, usually reaching its minimum at the first

1. This formula is valid for two’s complement number representation.

comb stage. However, if the stopband attenuation of the final lowpass filter *F(z)* is small, there may be slight truncation(s) between the comb stages, too.

Wordlength profiles shown in [2][3][4] resemble that of Figure 5, but they do not exploit the negative $\hat{G}_S$ ‘s.

5 CONCLUSIONS

In this paper, we have shown that in a digital signal processing system – contrary to the common knowledge – reducing the wordlength of a signal may sometimes reduce the total truncation noise variance at system output if noise shaping occurs. This phenomenon was deduced from the truncation noise model based on discrete error distribution and has been verified with simulations. It has to be taken into account in wordlength optimization algorithms, and may also help in manual selection of wordlengths.

If noise shaping – explicit or implicit – occurs in a wordlength-optimized DSP system, caution should be exercised if implementing wordlengths longer than those obtained from the optimization, as that may result in exceeding the given truncation noise constraints at system output.

It was also shown that CIC decimators provide inherent noise shaping, which helps reduce the complexity of both the CIC decimator itself and a successive filter stage.

REFERENCES

[1] G.A. Constantinides, P.Y.K. Cheung and W. Luk, Truncation Noise in Fixed-Point SFGs, Electronics Letters, p. 2012–2014, Vol. 35, No. 23, November 1999.

[2] E. Hogenauer, An Economical Class of Digital Filters for Decimation and Interpolation, IEEE Trans. Acoustics, Speech and Signal Proc., pp. 155–162, Vol. 29, No. 2, April 1981.

[3] S. Chu and C. S. Burrus, Multirate Filter Design Using Comb Filters, IEEE Trans. Circuits and Systems, pp. 913–924, Vol. 31, No. 11, November 1984.

[4] K.-Y. Khoo, Z. Yu, and A. N. Willson, Jr., Efficient High-Speed CIC Decimation Filter, Proc. IEEE Int. ASIC Conf. 1998, pp. 251–254.