14th European Signal Processing Conference (EUSIPCO 2006), Florence, Italy, September 4-8, 2006, copyright by EURASIP

DESIGN TRADE-OFFS FOR LINEAR-PHASE FIR DECIMATION FILTERS AND ΣΔ-MODULATORS

Anton Blad, Per Löwenborg, and Håkan Johansson

Electronics Systems, Linköping University
SE-581 83 Linköping, Sweden
phone: +46 13 281676, fax: +46 13 139282
email: {antonb, perl, hakanj}@isy.liu.se
web: http://www.es.isy.liu.se

ABSTRACT

In this paper we examine the relation between signal-to-noise-ratio, oversampling ratio, transition bandwidth, and filter order for some commonly used sigma-delta-modulators and corresponding decimation filters. The decimation filters are equi-ripple finite impulse response filters and it is demonstrated that, for any given filter order, there exists an optimum choice of the stopband ripple and stopband edge which minimizes the signal-to-noise-ratio degradation.

1. INTRODUCTION

The sigma-delta modulator (ΣΔ-modulator) is today often the preferred architecture for realizing low- to medium-speed analog-to-digital converters (ADCs) with effective resolution above 12 bits. Higher resolution than this is difficult to achieve for non-oversampled ADCs without laser trimming or digital error correction, since device matching-errors of semiconductor processes limit the accuracy of critical analog components [1]. The ΣΔ-modulator can overcome this problem by combining the speed advantage of analog circuits with the robustness and accuracy of digital circuits. Through oversampling and noise shaping, the ΣΔ-modulator converts precise signal waveforms to an oversampled digital sequence where the information is localized in a narrow frequency band in which the quantization noise is heavily suppressed.

The price to pay for these advantages is the required digital decimators operating on high sample-rate data. For Nyquist-rate CMOS ADCs, the power consumption increases approximately by a factor of four when increasing the resolution by one bit [2]. Hence, the power consumption of accurate Nyquist-rate ADCs tends to become very high. On the other hand, the analog circuitry of oversampled ΣΔ-modulators, does not need to be accurate and power savings can therefore be made, in particular for continuous-time input ΣΔ-modulators [3], provided that the power consumption of the necessary digital decimation filter is not too large. The digital filter is however a major source of power consumption of a high-resolution ΣΔ-modulator and needs to be carefully designed and optimized.

For FIR filter-based decimators, the filter complexity is to the largest extent determined by the transition bandwidth [4]. Further, a significant part of the quantization noise energy of an oversampled ΣΔ-modulator is located in this transition band, in particular when higher-order noise shaping is applied. Therefore, there exists a pronounced trade-off between decimator filter complexity and ΣΔ-modulator signal-to-noise-ratio (SNR). This paper presents investigations on how the SNR is degraded as a function of the filter transition bandwidth and filter order for various commonly used ΣΔ-modulator architectures and oversampling ratio (OSR). It is also demonstrated that, for a given filter order, there exists an optimum choice of the stopband ripple and stopband edge for equi-ripple filter solutions which minimizes the SNR degradation. Although ΣΔ-modulators have been known and used for quite some time, it appears that a thorough investigation of the relations just mentioned has not been published before.

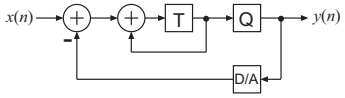

Figure 1: A first-order, single-feedback ΣΔ-modulator.

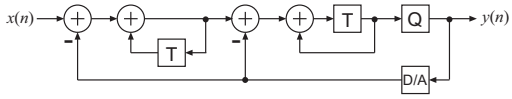

Figure 2: A second-order, double-feedback ΣΔ-modulator.

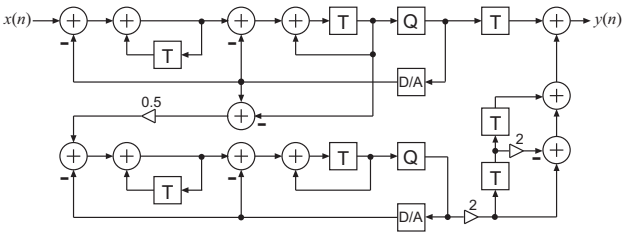

Figure 3: A multistage noise-shaping (MASH) ΣΔ-modulator.

2. SIGMA-DELTA-MODULATORS AND DECIMATION FILTERS

In the linear model of a ΣΔ-modulator structure, the input $x(n)$ and quantization error $e(n)$ are assumed uncorrelated. For analysis purposes, one then defines a signal transfer function $H(z)$ and noise transfer function $G(z)$ by which one obtains

$$Y(z) = H(z)X(z) + G(z)E(z)$$

where $Y(z)$, $X(z)$, and $E(z)$ are the z-transforms of the output $y(n)$, input $x(n)$, and quantization noise $e(n)$, respectively.

For the three topologies studied in this paper, shown in Figs. 1–3, viz, a first-order, single-feedback ΣΔ-modulator,

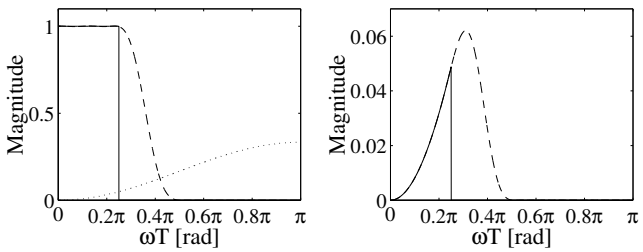

Figure 4: Illustration of output noise spectral density function for a 1-bit first-order sigma-delta using an ideal filter (solid line) and practical filter with a transition band (dashed line). The dotted line shows the spectral density function before filtering [see (1)].

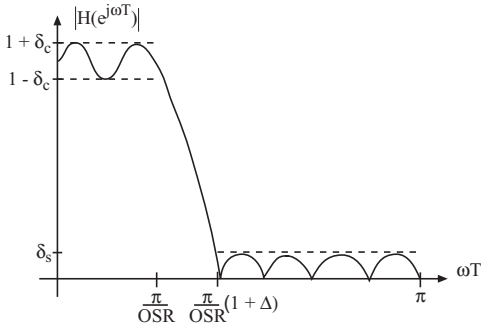

Figure 5: Decimation filter specification.

a second-order, double-feedback ΣΔ-modulator, and a multistage noise-shaping (MASH) ΣΔ-modulator, the signal transfer functions are pure delays according to

$$H_1(z) = z^{-1}$$
$$H_2(z) = H_{\mathrm{MASH}}(z) = z^{-2}$$

whereas the noise transfer functions have highpass characteristics according to

$$G_1(z) = (1-z^{-1})$$
$$G_2(z) = (1-z^{-1})^2$$
$$G_{\mathrm{MASH}}(z) = 2(1-z^{-1})^4$$

The signal power at the output is thus the same as that of the input whereas the noise is highpass shaped. Assuming that the quantization error $e(n)$ can be modeled as a wide-sense stationary process, with spectral density function $R_{ee}(e^{j\omega T})$, the corresponding spectral density function at the output, say $R_{y_e y_e}(e^{j\omega T})$, becomes

$$\begin{aligned} R_{y_e y_e}(e^{j\omega T}) &= R_{ee}(e^{j\omega T})\left|G(e^{j\omega T})\right|^2 \\ &= R_{ee}(e^{j\omega T}) K_0 \sin^{2K}\left(\frac{\omega T}{2}\right) \end{aligned} \qquad (1)$$

where $K_0 = 4, 16,$ and 1024 for the three ΣΔ-modulators in Figs. 1–3 with $K = 1, 2,$ and 4, respectively.

It is common to regard $e(n)$ as white noise with average power $\sigma_e^2 = 2^{-2(B-1)}/12$, B being the number of bits. In this case, $R_{ee}(e^{j\omega T}) = \sigma_e^2$ and the shape of $R_{y_e y_e}(e^{j\omega T})$ becomes typically as shown in Fig. 4 for $K = 1$ and 1-bit quantization ($\sigma_e^2 = 1/12$). The principle of a sigma-delta converter is to utilize the noise shaping together with oversampling, as exemplified in Fig. 4. By assuming that the signal content lies in the band $[0, \omega_c T]$, and by using a filter to remove out-of-band noise in the band $[\omega_c T, \pi]$, a much higher SNR than that corresponding to the quantizer alone can be achieved. However, in practice one can only approximate the ideal lowpass filter. This means that there will be a transition band between, say $\omega_c T$ and $\omega_s T$, where we have little or no control over the filter. Due to the smooth highpass shape of the noise, this will have a large degrading effect upon the SNR as a large portion of the noise after filtering remains in the transition band. This is seen in Fig. 4. In this example, the SNR at the output of the practical filter is 6.55 dB lower than the ideal filter's output which corresponds to slightly more than a one-bit degradation in resolution.

In principle, it is possible to approximate the ideal filter as closely as desired by increasing the filter order, but beyond a certain order this becomes intolerable from the implementation feasibility and cost points of view. It is therefore of interest to gain more insights into the relation between the SNR, OSR, which equals $\pi/\omega_c T$, transition bandwidth, and filter order. The main purpose of this paper is to provide such insights. Although ΣΔ-modulators have been known and used for quite some time, and many papers exist that deal with filter design and implementation for such modulators, it appears that a thorough investigation of the relations just mentioned cannot be found in the literature.

In our investigations, we will in this paper design the filters in the minimax sense. It may be argued that it would be more appropriate to instead use a least-squares design technique which aims at minimizing the energy instead of maximizing the stopband attenuation. This is indeed the natural choice when maximizing the SNR. However, there are usually additional restrictions that rather makes minimax design more appropriate than least-squares design. For example, in communication systems, it is common that blockers have to be suppressed by a certain amount. To control the level of suppression, it is necessary to incorporate additional "minimax constraints". This is why it is interesting to consider minimax design.

Finally, we like to make the following remark. In order to reduce the overall implementation complexity, it is common to make use of multistage decimators [5]. The requirements of each individual stage are thereby reduced with leads to a lowered overall cost. In particular, one can include recursive comb decimator structures which require very few arithmetic operations. It is stressed that our results by no means contradict the fact that such decimators are efficient, as our investigations are independent of the way the filtering and subsequent decimation are performed, i.e., they are structure independent. In other words, we study the relation between the stringency of the filter specification and the other parameters of interest. As a measure of the specification stringency we use the filter order because, for obvious reasons, this parameter increases (decreases) with increasing (decreasing) stringency. The results to be presented are thus valid regardless of the decimator structures that one wishes to use, including in particular the ones that utilize efficient recursive comb decimator stages.

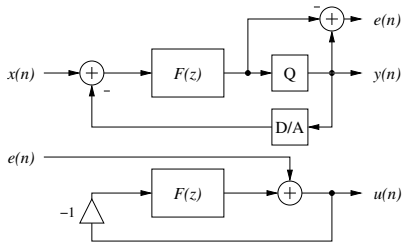

Figure 6: Principle of the extraction of quantization noise from a ΣΔ-modulator and the filtering thereof through a time-domain realization of the ΣΔ-modulator noise transfer function.

3. DESIGN CONSIDERATIONS

3.1 Decimation filter specification

In this paper, we use Nth-order symmetric linear-phase FIR filters that satisfy the specification shown in Fig. 5 where δ_c, δ_s, $\omega_c T = \pi/\text{OSR}$, and $\omega_s T = \pi(1+\Delta)/\text{OSR}$, denote the passband ripple, stopband ripple, passband edge and stopband edge, respectively. The filters are designed in the minimax sense using the well-known McClellan-Parks-Rabiner's (MPR) algorithm that is implemented in Matlab's function *firpm.m*. In this paper, the passband ripple is specified to be 0.01, but after each design it may be slightly smaller as there generally is a design margin. This is because the MPR algorithm only handles the ratio between the passband and stopband ripples.

The stopband edge is related to the passband edge through the relative transition bandwidth Δ that we define as [see Fig. 5]

$$\Delta = \frac{\omega_s T - \omega_c T}{\omega_c T}$$

3.2 Signal-to-noise-ratio

We denote the output signal and noise power (after the filtering) of the ΣΔ-modulator as P_{yx} and P_{ye}, respectively. Assuming that the input signal $x_a(t)$ is bandlimited to ω_c, and the passband ripple δ_c of the filter $H(z)$ is small enough so its effect on the signal power can be neglected, P_{yx} equals the input signal power, i.e., $P_{yx} = P_x$. The contribution to the output noise power emanates from the passband, transition band, and stopband regions. We denote the corresponding noise powers as $P_{ye}^{(pb)}$, $P_{ye}^{(tb)}$, and $P_{ye}^{(sb)}$, respectively. The total noise power is thus $P_{ye} = P_{ye}^{(pb)} + P_{ye}^{(tb)} + P_{ye}^{(sb)}$. The SNR at the output is given by

$$SNR = 10\log_{10}\frac{P_{yx}}{P_{ye}}$$

Further, we define the SNR degradation ΔSNR as the degradation in SNR caused by using a practical filter instead of an ideal lowpass filter. That is,

$$\Delta SNR = 10\log_{10}\frac{P_{yx}}{P_{ye}^{(pb)}} - 10\log_{10}\frac{P_{yx}}{P_{ye}} = 10\log_{10}\frac{P_{ye}}{P_{ye}^{(pb)}}$$

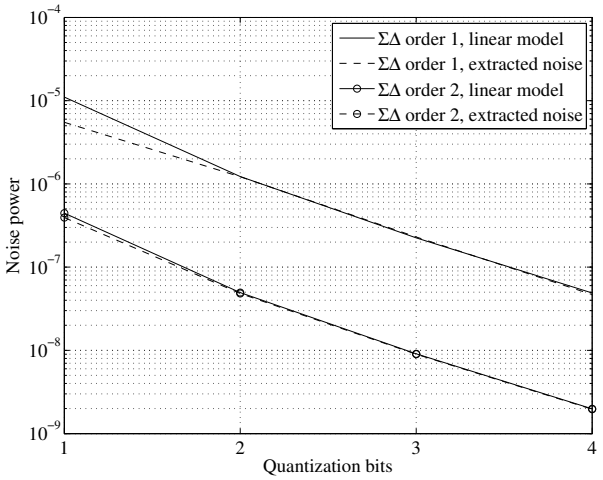

Figure 7: Quantization noise power of ΣΔ-modulators of order one and two, after decimation filter, according to the linear noise model and the computed quantization noise. The OSR is 32.

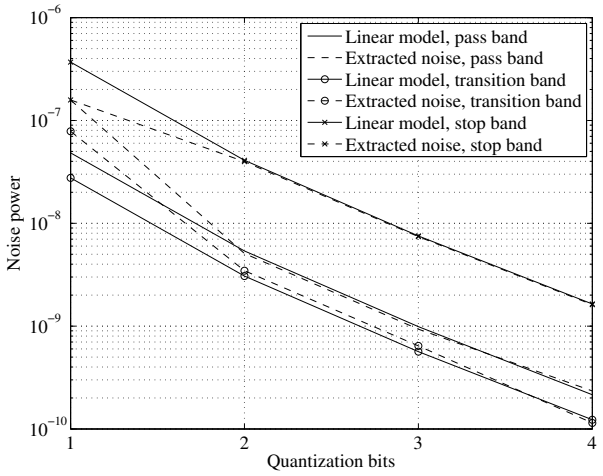

Figure 8: Quantization noise power of a second-order ΣΔ-modulator, in different bands of the decimation filter for the linear noise model and the computed quantization noise. The OSR is 32.

Using the linear model discussed in Section 2, the noise power is easily computed as

$$P_e = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{y_e y_e}(e^{j\omega T})\,\mathrm{d}(\omega T)$$

with $R_{y_e y_e}(e^{j\omega T})$ from (1). However, this model tends to be less appropriate for coarse quantization steps. In particular, problems arise for one-bit quantization. In this work, we have therefore evaluated the noise power through the model in Fig. 6, where the extracted quantization error $e(n)$ has been filtered through a time-domain realization of the ΣΔ-modulator noise transfer function $G(z)$. The so obtained sequence is then filtered through the decimation filter in order to get the final output noise. In Figs. 7 and 8, the noise power of the two different cases is compared.

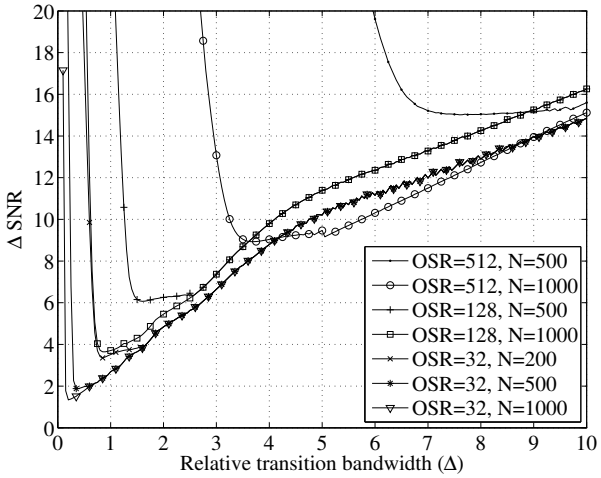

Figure 9: SNR degradation as a function of relative transition bandwidth Δ for the first-order ΣΔ-modulator.

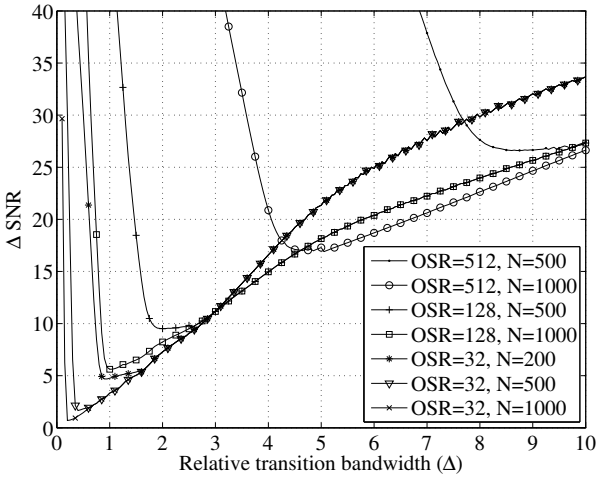

Figure 10: SNR degradation as a function of relative transition bandwidth Δ for the second-order ΣΔ-modulator.

4. SIMULATION RESULTS

For the three ΣΔ-modulator topologies shown in Figs. 1–3 we have performed two different investigations. In the first, the transition bandwidth of the decimation filter has been varied for different oversampling ratios and filter orders and the corresponding SNR degradation have been calculated. The results of this investigation are plotted in Figs. 9–11. For a given transition bandwidth and filter order the stopband ripple is fixed. From the figures, it can be seen that the SNR degradation has a minimum for a certain choice of Δ.

In the second investigation, the optimal choice of the relative transition bandwidth Δ has been found for filter orders between 100 and 1000 for different oversampling ratios. The optimal Δ's for the three ΣΔ-modulator topologies are shown in Figs. 12–14. The corresponding SNR degradations can be found in Figs. 15–17. Decreasing the transition bandwidth below the optimum causes the SNR to worsen considerably, because of rapidly decreasing stopband attenuation. Also, for large enough transition bands, increasing the filter order will yield no significant SNR improvements because essentially all the noise power is from the transition band region. It should be noted that optimal SNR may occur for a transition bandwidth several times wider than the passband width.

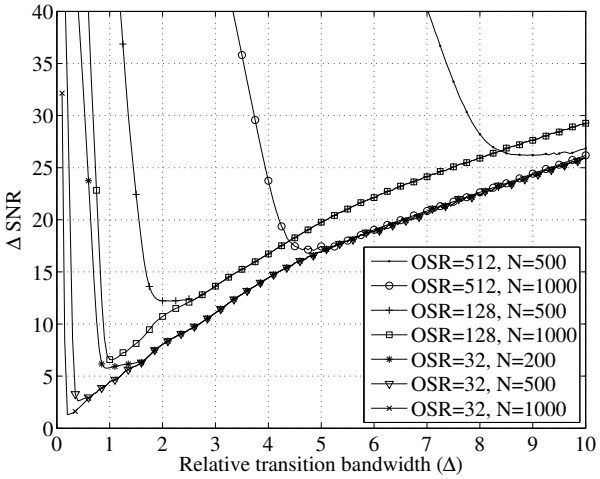

Figure 11: SNR degradation as a function of relative transition bandwidth Δ for the MASH ΣΔ-modulator.

5. CONCLUSION

In this paper, design trade-offs for linear-phase FIR decimation filters and ΣΔ-modulators have been investigated. The results are useful for designers of ΣΔ-modulator-based ADCs in finding a balanced overall solution that sets reasonable requirements on the modulator as well as the FIR decimation filter.

REFERENCES

[1] R. Schreier and G. C. Temes, *Understanding Delta-Sigma Data Converters*, New Jersey: Wiley, 2005.

[2] K. Uyttenhove and M. Steyaert, "Speed-power-accuracy trade-off in high-speed ADC's," *IEEE Trans. Circuits Syst. II*, vol. 4, pp. 247-257, Apr. 2002.

[3] P. G. A. Jespers, *Integrated Converters: D to A and A to D Architectures, Analysis and Simulation*, New York: Oxford University Press, 2001.

[4] O. Herrmann, R. L. Rabiner, and D. S. K. Chan, "Practical design rules for optimum finite impulse response digital filters," *Bell Tech. J.*, vol. 52, no. 6, pp. 769-799, 1973.

[5] J. C. Candy and G. C. Temes, "Oversampling methods for A/D and D/A conversion", in *Oversampling Delta-Sigma Data Converters*, Eds. J. C. Candy and G. C. Temes, IEEE Press, 1992.

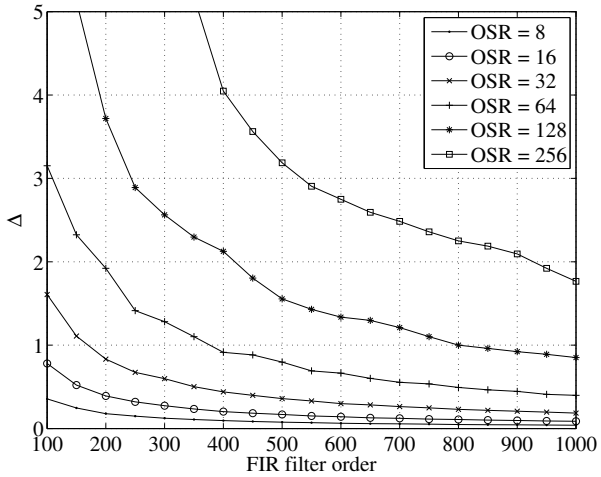

Figure 12: SNR-optimal choice of Δ as a function of the filter order for the first-order ΣΔ-modulator.

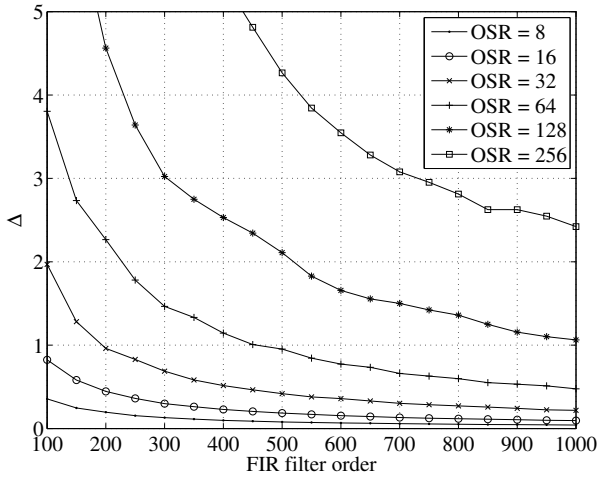

Figure 13: SNR-optimal choice of Δ as a function of filter order for the second-order ΣΔ-modulator.

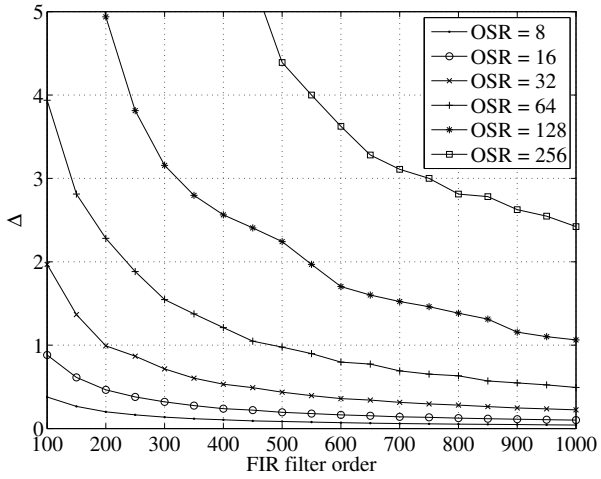

Figure 14: SNR-optimal choice of Δ as a function of filter order for the MASH ΣΔ-modulator.

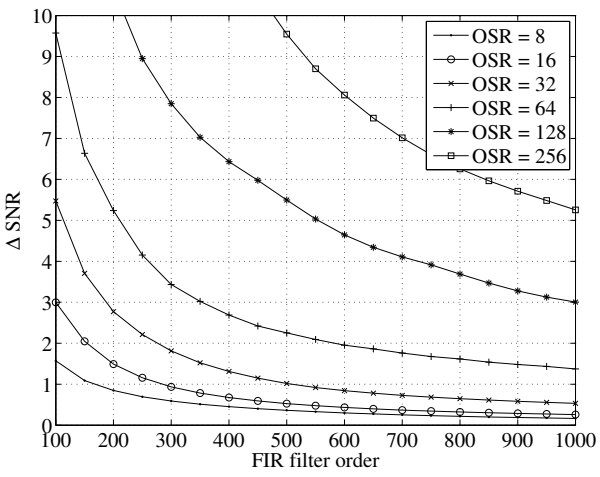

Figure 15: Minimal SNR degradation as a function of filter order for the first-order ΣΔ-modulator.

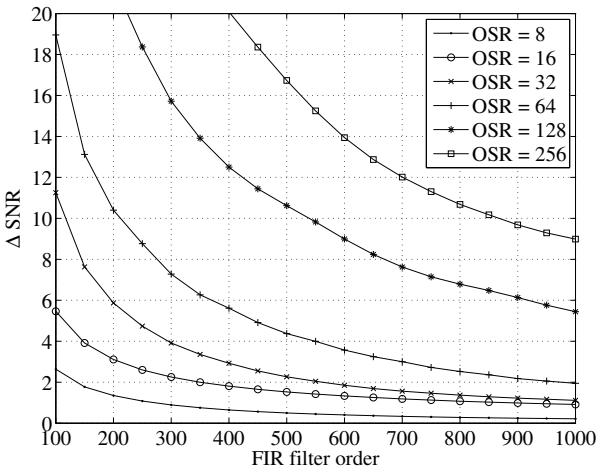

Figure 16: Minimal SNR degradation as a function of filter order for the second-order ΣΔ-modulator.

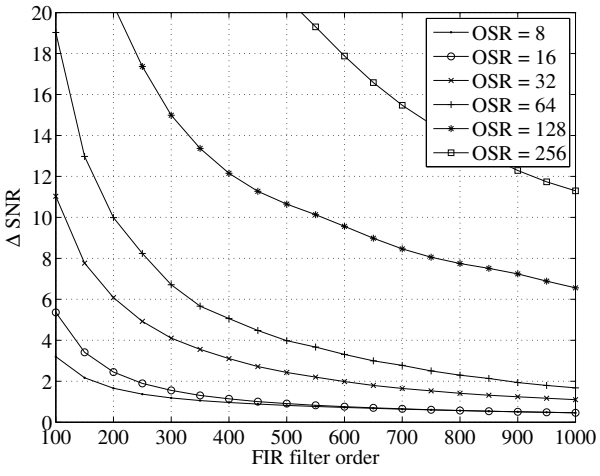

Figure 17: Minimal SNR degradation as a function of filter order for the MASH ΣΔ-modulator.