

Chapter 2

Basics of Sigma-Delta Modulation

The principle of sigma-delta modulation, although widely used nowadays, was developed over a time span of more than 25 years. Initially the concept of oversampling and noise shaping was not known and the search for an efficient technique for transmitting voice signals digitally resulted in the Delta Modulator. Delta modulation was independently invented at the ITT Laboratories by Deloraine et al. [11, 12] the Philips Research Laboratories by de Jager [10], and at Bell Telephone Labs [8] by Cutler. In 1954 the concept of oversampling and noise shaping was introduced and patented by Cutler [9]. His objective was not to reduce the data rate of the signal to transmit as in earlier published work, but to achieve a higher signal-to-noise ratio in a limited frequency band. All the elements of modern sigma-delta modulation are present in his invention, except for the digital decimation filter required for obtaining a Nyquist rate signal. The name Delta-Sigma Modulator (DSM) was finally introduced in 1962 by Inose et al. [25, 26] in their papers discussing 1-bit converters. By 1969 the realization of a digital decimation filter was feasible and described in a publication by Goodman [16]. In 1974 Candy published the first complete multi-bit Sigma-Delta Modulator (SDM) in [6]. Around the same time the name SDM was introduced as an alternative for Delta-Sigma Modulator and since then both names are in use. In this book the oversampled noise-shaping structure will be referred to as SDM. According to the author SDM is the more appropriate name since the integration or summing (the sigma) is over the difference (the delta).

In the 70's, because of the initially limited performance of Sigma-Delta Modulators, their main use was in encoding low frequency audio signals (analog-to-digital conversion) using a 1-bit quantizer and a first or a second order loop filter. The creation of black and white images for print from a gray scale input was another application where Sigma-Delta noise-shaping techniques were used (digital-to-digital conversion). Since then a lot of research on improving SDM performance has been performed and great improvements have been realized. Nowadays top of the line SDM based analog-to-digital converters (ADCs) use a multi-bit quantizer and a high-order loop filter and are capable of converting 10's of MHz of bandwidth with high dynamic range. Because of high power efficiency, Sigma-Delta based analog-to-digital converters are used in the radio of mobile telephones. Another example




Fig. 2.1 Oversampling does not affect the signal power or total quantization noise power but reduces the noise spectral density

of the efficient use of sigma-delta modulation techniques is the Super Audio CD format which uses a 64 times oversampled 1-bit signal for delivering a 120 dB signal-to-noise ratio (SNR) over the 0–20 kHz band. In this specific example the decimation filter is omitted and the oversampled signal is directly stored as to minimize signal operations and therefore maximize the signal quality. An omnipresent example of sigma-delta modulation in digital-to-analog conversion can be found in portable audio playback devices, e.g. IPOD and MP3 players. The audio digital-to-analog converter (DAC) in these devices realizes its performance using noise shaping (NS) and pulse-width-modulation (PWM) or pulse-density-modulation (PDM) techniques. These PWM/PDM signals are typically generated using a (modified) digital SDM.

Although all these SDM solutions are optimized for a certain application and context, they still share the same underlying basic principles of oversampling and noise shaping. Oversampling is the process of taking more samples per second than required on the basis of the Nyquist-Shannon criterion. By changing the sampling rate the signal power and total quantization noise power is not affected. Therefore, the signal to quantization noise ratio is not changed. However, the quantization noise is spread over a larger frequency range, reducing the spectral density of the quantization noise. If now only the original Nyquist band is considered, the quantization noise power is reduced by 3 dB for every doubling of the oversampling ratio and the signal to quantization noise ratio is improved accordingly. This effect is illustrated in Fig. 2.1 for an oversampling ratio (OSR) of 1, 2, and 4 times.

Noise shaping is applied as a second step to improve the signal to quantization noise ratio. In this process the frequency distribution of the quantization noise is altered such that the quantization noise density reduces in the signal band. As a result the noise density increases at other frequencies where the noise is less harmful. This effect is depicted in Fig. 2.2, where low frequency noise is pushed to high frequencies. The amount of quantization noise is not changed by this process but the signal-to-noise ratio is increased in the low frequency area of the spectrum. In an SDM the techniques of oversampling and noise shaping are combined, resulting in

Fig. 2.2 Low frequency noise is pushed to high frequencies by noise shaping





Fig. 2.3 Generic model of the Sigma-Delta noise-shaping loop, consisting of 2-input loop filter and quantizer



an increased efficiency since now the quantization noise can be pushed to frequencies far from the signal band.

All SDM structures realize the shaping of noise with an error minimizing feedback loop in which the input signal x is compared with the quantized output signal y , as depicted in Fig. 2.3. The difference between these two signals is frequency weighed with the loop filter. Differences between the input and output that fall in the signal band are passed to the output without attenuation, out-of-band differences are suppressed by the filter. The result of the weighing is passed to the quantizer, which generates the next output value y . The output y is also fed back to the input, to be used in the next comparison. The result of this strategy is a close match of input signal and quantized output in the pass-band of the filter, and shaping of the quantization errors such that those fall outside the signal band.

In Sect. 2.1 the noise-shaping loop in data converters will be examined in detail, revealing that in reality only analog-to-digital (AD) and digital-to-digital (DD) noise shaping conversion exists. Over the last decennia a great variety of noise-shaping loops have been developed, but all originate from a minimal number of fundamental approaches. The most commonly used configurations are discussed in Sect. 2.2. During the design phase of an SDM the noise-shaping transfer function is typically evaluated using a linear model. In reality, especially for a 1-bit quantizer, the noise transfer is highly non-linear and large differences between predicted and actual realized transfer can occur. In Sect. 2.3 the linear modeling of an SDM is examined and it will be shown that simulations instead of calculations are required for evaluating SDM performance. Several criteria exist for evaluating the performance of an SDM. The criteria can be differentiated between those that are generic and are used for characterizing data converters in general, and those that are only applicable for Sigma-Delta converters. Both types are discussed in Sect. 2.4.

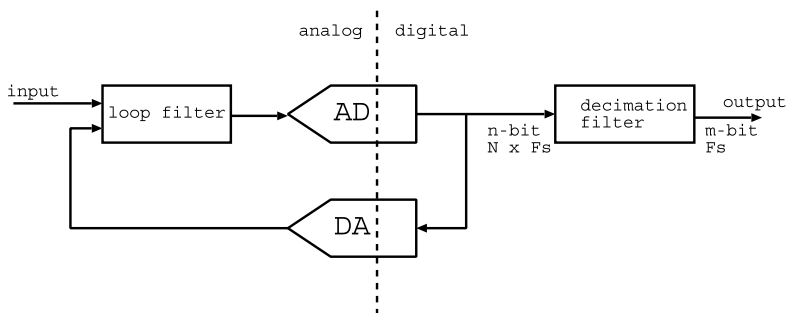


Fig. 2.4 Main building blocks of a Sigma-Delta analog-to-digital converter

2.1 AD, DD, and DA Sigma-Delta Conversion

2.1.1 AD Conversion

The most well-known form of sigma-delta modulation is analog-to-digital conversion. In Fig. 2.4 the main building blocks of a generic Sigma-Delta ADC are shown. In the figure the analog and digital domains are indicated as well. The analog signal that will be converted, as well as the DAC feed-back signal, enter the analog loop filter at the left side of the figure. The output of the loop filter is converted to an n -bit digital signal by the quantizer (ADC). This n -bit digital signal is passed to a digital decimation filter and to the feed-back DAC. The decimation filter removes the out-of-band quantization noise, thereby converting the high rate low resolution signal to a high resolution low rate signal. The feed-back DAC performs the inverse function of the ADC (quantizer) and converts the n -bit digital code to an analog voltage or current, closing the Sigma-Delta loop.

Several different types of analog Sigma-Delta Modulators exist, varying in for example the way the loop filter is functioning (e.g. continuous time or discrete time) or how the DAC is constructed (e.g. switched capacitor or resistor based). Independent of these details, in all structures the use of a low resolution ADC and DAC is key. The coarse quantization results in a large amount of quantization noise which is pushed out of band by the loop filter. The number of bits used in the ADC and DAC is typically in the range 1–5. A 1-bit quantizer is easier to build than a 5-bit quantizer, requires less area and power, and is intrinsically linear, but has the disadvantage that less efficient noise shaping can be realized and that a higher oversampling ratio is required to compensate for this. The final Sigma-Delta output, i.e. at the output of the decimation filter, will be an m -bit word where m can be as high as 24. The number of bits is independent on the number of bits used in the internal ADC and DAC. Sometimes only the part before the decimation filter is considered in discussions about Sigma-Delta Modulators.




Fig. 2.5 Main building blocks of a Sigma-Delta digital-to-digital converter

2.1.2 DD Conversion

In a digital-to-digital Sigma-Delta converter an n -bit digital input is converted to an m -bit digital output, where n is larger than m . The sampling rate of the signal is increased during this process in order to generate additional spectral space for the quantization noise. The main building blocks of a generic DD SDM are shown in Fig. 2.5. The n -bit signal is first upsampled from F_s to $N \times F_s$ in the upsampling filter. The resulting signal is passed to the actual SDM loop. This loop is very similar to the one in Fig. 2.4, except that now everything is in the digital domain. The ADC and DAC combination is replaced by a single quantizer which takes the many-bit loop-filter output and generates a lower-bit word. Since everything is operating in the digital domain no DAC is required and the m -bit word can directly be used as feed-back value. The noise-shaped m -bit signal is the final Sigma-Delta output. This m -bit signal is often passed to a DA converter, resulting in a Sigma-Delta DAC. In the case of audio encoding for Super Audio CD the 1-bit output is the final goal of the processing and is directly recorded on disc.

2.1.3 DA Conversion

A Sigma-Delta based DA converter realizes a high SNR with the use of a DAC with few quantization levels and noise-shaping techniques. In the digital domain the input signal to the DAC is shaped, such that the quantization noise of the DAC is moved to high frequencies. In the analog domain a passive low-pass filter removes the quantization noise, resulting in a clean baseband signal. The structure of a Sigma-Delta DAC is, except for some special PWM systems, a feed-forward solution, i.e. there is no feed-back from the analog output into the noise-shaping filter. Because the noise-shaping feed-back signal is not crossing the analog-digital boundary, the name Sigma-Delta DAC is confusing and misleading. A Sigma-Delta DAC is the combination of a DD converter and a high-speed few-bit DAC. In Fig. 2.6 the complete Sigma-Delta DAC structure is shown. The digital n -bit input signal is passed to a DD converter which upsamples the input to $N \times F_s$ before an all digital SDM reduces the word-length. The noise-shaped m -bit signal is passed to the m -bit DAC which converts the digital signal to the analog domain. Finally the analog signal is filtered to remove the out-of-band quantization noise.




Fig. 2.6 Main building blocks of a Sigma-Delta digital-to-analog converter


Fig. 2.7 Generic model of the Sigma-Delta noise-shaping loop, consisting of 2-input loop filter and quantizer



2.2 Sigma-Delta Structures

In Sect. 2.1 it was shown that two basic SDM types exist, i.e. with an analog or a digital loop filter. In the case of an analog filter the combination of a quantizing ADC and a DAC is required for closing the noise-shaping loop and a decimation filter is present at the output. In the case of a digital filter no analog-digital domain boundary has to be crossed and only a digital quantizer is required, but at the input an upsample filter is present. When studying the noise-shaping properties of an SDM from a high-level perspective these analog-digital differences can be safely ignored and a generic model of the Sigma-Delta noise-shaping loop can be used instead. This generic model, consisting of a loop filter and a quantizer, is depicted in Fig. 2.7. The loop filter has two inputs, one for the input signal and one for the quantizer feed-back signal, where the transfer function for the two inputs can be complete independent in theory. In practice large parts of the loop-filter hardware will be shared between the two inputs. A practical loop-filter realization will consist of addition points, integrator sections, feed-forward coefficients b_i and feed-back coefficients a_i as shown in Fig. 2.8. In this structure the number of integrator sections sets the filter order, e.g. 5 concatenated integrators results in a fifth order filter. The exact filter transfer is realized by the coefficients. With proper choice of b_i and a_i the complexity of the filter structure can be reduced, e.g. resulting in a feed-forward structure. This optimized structure can be redrawn to give a 1-input loop filter where the first subtraction is shifted outside the filter, as depicted in Fig. 2.9. As an alternative it is possible make all b_i equal to zero except for b_N and realize the noise-shaping transfer using only a_i . This structure is referred to as a feed-back SDM and is shown in Fig. 2.10. The two structures can be made to behave identical in terms of noise shaping but will realize a different signal transfer. In both structures the quantizer can have any number of quantization levels. In practice values between 1-bit (2 levels) and 5-bit (32 levels) are used.

Fig. 2.8 Internal structure of practical 2-input loop filter, consisting of integrators, subtraction points, feed-forward coefficients b_i and feed-back coefficients a_i



As an alternative to the single-loop SDM with multi-bit quantizer, a cascade of first-order Sigma-Delta Modulators can be used. This structure is commonly referred to as multi-stage noise shaping (MASH) structure. In an MASH structure the quantization error of a first modulator is converted by a second converter, as depicted in Fig. 2.11. By proper weighing the two results in the digital domain with filters $H1$ and $H2$ the quantization noise of the first modulator is exactly canceled




Fig. 2.9 SDM with feed-forward loop filter. The subtraction point of signal and feedback has been shifted outside the loop filter




Fig. 2.10 SDM with feedback loop filter





Fig. 2.11 Second order MASH SDM

Fig. 2.12 Noise shaper structure



and only the shaped noise of the second modulator remains. In this fashion an n th order noise shaping result can be obtained by using only first order converters. The disadvantage compared to a single-loop SDM is the inability to produce a 1-bit output.

Closely related to the SDM is the noise shaper structure. In a noise shaper no filter is present in the signal path and only the quantization error is shaped. This is realized by inserting a filter in the feed-back path which operates on the difference between the quantizer input and quantizer output, as depicted in Fig. 2.12. With a proper choice of the filter the same noise shaping can be realized as with an SDM. Unique for the noise shaper is that only the error signal is shaped and that the input signal is not filtered. Because of this special property the noise shaper can also be used on non-oversampled signals to perform in-band noise shaping. This technique is, for example, used to perform perceptually shaped word-length reduction for audio signals, where 20-bit pulse-code modulated (PCM) signals are reduced to 16-bit signals with a higher SNR in the most critical frequency bands at the cost of an increase of noise in other frequency regions.

2.3 Linear Modeling of an SDM

For a generic discrete-time SDM in feed-forward configuration, as depicted in Fig. 2.13, the signal transfer function (STF) and noise transfer function (NTF) will be derived on the basis of a linear model. In this figure $x(k)$ represents the discrete-time input signal, $d(k)$ the difference between the input and the feed-back signal (the instantaneous error signal), $H(z)$ is the loop filter, $w(k)$ the output of the loop filter (the frequency weighted error signal), and $y(k)$ is the output signal.

Fig. 2.13 Generic model of a digital SDM in feed-forward configuration





Fig. 2.14 Linear model of a digital SDM in feed-forward configuration



The difference between the quantizer output $y(k)$ and quantizer input $w(k)$ is the quantization error $e(k)$. For the schematic we can write:

$$\begin{aligned} y(k) &= w(k) + e(k) \\ &= H(z) \cdot [x(k) - y(k)] + e(k) \end{aligned} \quad (2.1)$$

$$y(k) \cdot [1 + H(z)] = H(z) \cdot x(k) + e(k) \quad (2.2)$$

$$y(k) = \frac{H(z)}{1 + H(z)} \cdot x(k) + \frac{1}{1 + H(z)} \cdot e(k) \quad (2.3)$$

From Eq. 2.3 it can be seen that the output signal $y(k)$ consists of the sum of a filtered version of the input $x(k)$ and a filtered version of the quantization error $e(k)$.

If it is assumed that the quantization error is not correlated with the input signal, the quantizer can be modeled as a linear gain g and an additive independent noise source $n(k)$ which adds quantization noise. The resulting linear SDM model is depicted in Fig. 2.14.

By replacing $e(k)$ in Eq. 2.3 with $n(k)$ and moving gain g into filter $H(z)$, the output $y(k)$ can now be described as

$$y(k) = \frac{H(z)}{1 + H(z)} \cdot x(k) + \frac{1}{1 + H(z)} \cdot n(k) \quad (2.4)$$

By setting $n(k) = 0$ the signal transfer function (STF) is obtained:

$$STF_{FF}(z) = \frac{y(k)}{x(k)} = \frac{H(z)}{1 + H(z)} \quad (2.5)$$


The signal transfer function is specific for the feed-forward structure, indicated by the subscript FF.

The noise transfer function (NTF) describes how the quantization noise, which is introduced by the quantization operation, is transferred to the output of the modulator. It is obtained by setting $x(k) = 0$ in Eq. 2.4:

$$NTF(z) = \frac{y(k)}{n(k)} = \frac{1}{1 + H(z)} \quad (2.6)$$

In order to realize a high signal-to-noise ratio in the baseband, the quantization noise should be suppressed for low frequencies and shifted to high frequencies.

Fig. 2.15 Transfer of a typical fifth order loop filter designed according to a Butterworth specification with 100 kHz corner frequency and additional resonator sections at 12 and 20 kHz. The sampling rate is 2.8 MHz



As a result the loop filter $H(z)$ should be a filter that provides a lot of gain for low frequencies and little gain for high frequencies, i.e. a low-pass characteristic. With $H(z)$ low-pass it can be appreciated that the STF will be close to unity for low-frequencies and that the input signal will be accurately captured. The transfer characteristic of a typical fifth order loop filter is plotted in Fig. 2.15. In this example the loop filter is designed according to a Butterworth specification for a corner frequency of 100 kHz when the sampling rate is 2.8 MHz (a 64 times oversampled 44 100 Hz system). Resonators (linear feed-back within the loop filter) at 12 and 20 kHz have been added for increasing the SNR [7, 50].

With $H(z)$ given, the linearized STF and NTF can be plotted using Eq. 2.5 and Eq. 2.6. The result for the STF for a feed-forward (FF) as well as a feed-back (FB) modulator is plotted in Fig. 2.16 for an assumed quantizer gain of 1.0. As expected, the STF equals unity for low frequencies for both types. Around the corner frequency of the feed-forward filter a gain of approximately 7 dB is realized before the filter starts to attenuate the input signal. At $F_s/2$ the input is attenuated by about 7 dB. The feed-back filter realizes a gain of approximately 3 dB at the corner frequency and then falls off strongly.

Plotting the NTF accurately is far less trivial. It has to be realized that Eq. 2.6 will only give a rudimentary approximation of the actual quantization noise spectrum, i.e. in Eq. 2.6 the quantization noise is treated as an independent signal whereas in reality the signal is depending on the quantizer input. Only if signal $e(k)$ is uncorrelated with the input signal, Eq. 2.6 will accurately describe the quantization noise. In the case of a multi-bit quantizer the quantization error is reasonably white for typical input signals. If desired, it can be made completely white by adding to the quantizer input a dither signal with triangular probability density (TPDF) that spans two quantization levels [54]. In the case of a single-bit quantizer the quantization error is strongly correlated with the input signal. Furthermore, since only two quantization levels exist it is not possible to add a TPDF dither signal of large enough amplitude to the quantizer input without overloading the modulator. In the case of a single-bit quantizer a deviation from the predicted NTF is therefore to be expected. Typical effects caused by the gross non-linearity of the 1-bit quantizer

Fig. 2.16 Linearized signal transfer function for the fifth order loop filter of Fig. 2.15 in feed-forward and feed-back configuration (quantizer gain of 1.0)





Fig. 2.17 Linearized noise transfer function of the fifth order loop filter of Fig. 2.15 (quantizer gain of 1.0)



are signal distortion, idle tones, and signal dependent baseband quantization noise (noise modulation).

In Fig. 2.17 the linearized NTF resulting from the 100 kHz filter is plotted for an assumed quantizer gain of 1.0. According to this prediction the quantization noise will be rising with 100 dB per decade and from 100 kHz onwards the spectrum will be completely flat. At 12 and 20 kHz a notch in the quantization noise floor should be present. By means of simulations the accuracy of this prediction will be verified. For a modulator with 1-bit quantizer the output spectrum for a 1 kHz input sine wave with an amplitude of -6 dB is plotted in Fig. 2.18 in combination with the predicted quantization noise spectrum. The FFT length used is 256 000 samples. The spectrum has been power averaged 16 times in order to obtain a smooth curve (see Appendix A). In the figure the predicted 100 dB per decade rise of the noise can be clearly identified. The high frequency part of the spectrum, however, deviates strongly from the prediction, i.e. a tilted noise floor with strong peaking close to $F_s/2$ is identified. In the baseband part of the output odd signal harmonics can be identified, which are not predicted by the linear models STF. The predicted notches

Fig. 2.18 Simulated output spectrum of an SDM with 1-bit quantizer and loop filter of Fig. 2.15. Input signal is a 1 kHz sine wave with an amplitude of -6 dB. The predicted quantization noise spectrum is indicated as a *dashed line*





Fig. 2.19 Simulated output spectrum of an SDM with 1-bit quantizer and loop filter of Fig. 2.15. Input signal is a DC of $1/128$. The predicted quantization noise spectrum is indicated as a *dashed line*



at 12 and 20 kHz are present. As a second example, for the same modulator the output spectrum for a DC input of $1/128$ is plotted in Fig. 2.19 and compared with the predicted quantization noise spectrum. The spectrum shows globally the same noise shaping as in the first example, with superimposed on it a large collection of discrete tones. These so-called idle tones cannot be understood from the linear model, but can clearly be an issue as they are not only present at high frequencies but also in the baseband.

As is clear from the two examples, large differences can exist between the prediction based on the linear model and actual modulator output in the case of a 1-bit quantizer. Since no accurate mathematical models for predicting a modulators response exist, the only reliable solution for obtaining performance figures of a 1-bit SDM is to perform time-domain simulations and analyze these results. Unfortunately, at the start of a design no realization exists yet and the linearized STF and NTF formulas have to be used for designing the initial loop filter. As a next step, computer simulations will have to be used to verify the response. Depending on the simulation outcome parameters will be iteratively adjusted until the desired result

is obtained. In order to obtain reproducible and comparable results, in this book the iterative approach for designing loop filters is not taken. Filters are designed using the linear model of a traditional SDM, according to a predetermined criterion, and used as-is. The predetermined criterion will typically be a transfer characteristic according to a Butterworth prototype filter with a specified corner frequency. The actual resulting transfer might be varying as a function of the input signal and the noise-shaping structure used, and can therefore only be compared by keeping the same filter which is designed using one and the same standard approach.

2.4 Sigma-Delta Modulator Performance Indicators

The performance of an SDM can be expressed in terms that describe the quality of the signal conversion process, as well as in terms of resources or implementation costs. The signal conversion performance can again be divided in two groups, namely performance measures that hold for data converters in general, and Sigma-Delta converter specific functional performance. The SDM specific functional performance indicators, discussed in Sect. 2.4.2, relate to the stability of the converter, limit cycle and idle tone behavior, noise modulation, and transient performance. In order to enable an easy comparison of designs, often a Figure-of-Merit (FoM) calculation is used. In the FoM several performance indicators are combined into a single number that represents the efficiency of a design. In the case of an SDM this approach is not straightforward, and the topic is therefore discussed separately.


2.4.1 *Generic Converter Performance*

The most often used generic data converter performance indicators are the Signal-to-Noise-Ratio (SNR), the Signal-to-Noise-and-Distortion-Ratio (SINAD or SNDR), the Spurious-Free-Dynamic-Range (SFDR), and Total-Harmonic-Distortion (THD). Next to these signal conversion performance metrics, the implementation and resource costs are important quality aspects of a converter. By combining several of these performance indicators into an FoM, the converter performance can be specified with a single value.

2.4.1.1 SNR and SINAD

The SNR and SINAD are two closely related measures. In both cases the harmonic signal power is compared to the power of the residual (noise) signal. The residual power can be split in noise and signal distortion components. In an SINAD measurement no differentiation between the two types of signal is made and the complete residual signal is integrated, hence the name Signal-to-Noise-and-Distortion-Ratio. In an ideal SNR measurement only the noise part of the residual signal is integrated.

Fig. 2.20 Example SDM output spectrum. FFT length is 256 000 samples, 16 power averages have been performed



In practice however, an SNR measurement will typically only ignore the harmonically related signal components. Non-harmonically related components, i.e. combinations of the input signal frequency and the clock frequency, are often treated as noise. The SNR figure is typically slightly higher than the SINAD value because of the absence of the harmonic components. Only in the case of no distortion the two numbers are equal.

In case of a Nyquist converter the noise integration is typically performed over the complete frequency band from 0 to $F_s/2$. In the case of an oversampled converter, e.g. an SDM, the integration is performed over the band of interest only. In this book the band of interest is the baseband part of the output, i.e. the frequency span of 0 to 20 kHz.

Since only part of the output spectrum is used for the SINAD calculations, the SINAD will typically show a strong input frequency dependency. Typical distortion of an SDM consists of odd harmonic components, i.e. components at $(2n + 1) \cdot f_{in}$. As an example, if the input frequency is chosen as 5 kHz, there will be harmonic components at 15 kHz, 25 kHz, 35 kHz, etc. Since only the baseband (0–20 kHz in most examples) is considered for SINAD calculations, only the component at 15 kHz will be taken into account. The SINAD value for this input frequency will therefore be most likely higher than for a slightly lower input frequency which has multiple harmonics in the baseband. In order to get a single representative SINAD number, i.e. one which takes most harmonic distortion components into account, in most experiments an input frequency of 1 kHz is used. For this frequency the first 19 harmonics fall within the signal band.

In Fig. 2.20 an example SDM output spectrum is shown. The input signal which has an amplitude of 0.5 (−6.02 dB) is visible at approximately 1 kHz (992 Hz), and harmonics at 3 kHz and 5 kHz are also clearly visible. In this example the SNR equals 113.2 dB and the SINAD equals 111.5 dB. The difference of 1.7 dB is primarily caused by the power in the third harmonic (HD3) and the fifth harmonic (HD5). Note that it is in general not possible to accurately read the SNR or SINAD value directly from a spectral plot — integration over all frequency bins is required and the spectral density per bin is a function of the number of points of the FFT. If

a large distortion component is present in the output a rough estimate of the SINAD can be made by subtracting the power of this component from the power of the fundamental.

2.4.1.2 SFDR

The SFDR is the difference in power between the test signal and the largest non-signal peak in the spectrum. The non-signal peak can be harmonically related but this is not required. In oversampled systems not the complete spectrum is taken into account, only the band of interest is considered. In the case of a digital SDM no artifacts other than those generated by the modulator itself are expected to be present, therefore typically the biggest peak is a harmonic component or the in-band rising noise-floor. In the example spectrum of Fig. 2.20 the third harmonic is the biggest non-signal component with a power of -123.4 dB, resulting in an SFDR of -6.0 dB $- (-123.4$ dB) = 117.2 dB.

2.4.1.3 THD

The THD is the ratio between the power in all the harmonic components and the signal power. In oversampled systems only the harmonic power in the band of interest is included in the calculation. The THD value relates to the linearity of a converter, i.e. a lower THD value means less signal dependent distortion. The THD is often a function of the input level. In analog converters large inputs typically cause circuits to saturate or clip and therefore generate distortion. In a digital SDM saturation and clipping can be avoided by using large enough word widths, but a 1-bit SDM will still generate harmonics, especially for large input signals. Determining the THD accurately can be difficult when the harmonic distortion components are of the same order of magnitude as the random noise components. In order to still get accurate results the technique of coherent averaging can be applied. The result of this process is that random frequency components are suppressed while coherent (signal) components are not. Every doubling of the number of averages reduces the random signals by 3 dB, e.g. performing 32 averages reduces the noise floor by 15 dB. Please refer to Appendix A for more details.

In the example of Fig. 2.20 the THD equals -116.3 dB, i.e. the combined power of all the harmonic components is 116.3 dB less than the power in the 1 kHz signal tone.

2.4.1.4 Implementation and Resource Costs

The costs of making a data converter fall in three main categories. First, there is the time required to design the converter. Second, there is the cost associated with the physical IC realization, i.e. materials and processing cost. Third, there is the

cost related to the industrial testing of the manufactured device. Next to these cost factors which are occurring only once, there is a reoccurring cost factor, i.e. the cost associated with the use of the converter. This cost manifests itself as the power consumption of the converter.

Both the silicon area and required design time depend on the type and the specifications of the converter, as well as on the experience of the designer. In general it holds that if the performance specification is more difficult to reach, the required design time will be longer and often the circuit will be bigger. The power consumption of the circuit typically also scales with the area and the performance level. For example, in AD converters often thermal noise is limiting the SNR. In order to increase the SNR, i.e. reduce the thermal noise, typically a larger current is required, which in turn requires larger active devices. A data converter that uses little power is preferred over a converter that requires a lot of power. A smaller silicon area results in less direct manufacturing cost. However, the industrial testing that is required can add significant cost. A converter that requires little testing is therefore preferred over a converter that requires a lot of tests.

2.4.1.5 Figure-of-Merit

Comparison of the power efficiency of two AD converters that achieve identical signal conversion specifications, i.e. have the same sampling rate and realize the same SNR for every input signal, is an easy task; the one with the lowest power consumption is the best. However, if the signal conversion specifications are not 100% identical, the comparison becomes difficult. To overcome this problem and make the comparison of different data converters possible, typically a Figure-of-Merit (FoM) is calculated. In the FoM a single value is used to represent the performance specifications of the converter, typically the power consumption and the signal conversion bandwidth and resolution.

Unfortunately, no universally agreed standard exists for calculation of the FoM. An often used FoM equation for the characterization of AD converters equals

$$FoM = \frac{P}{2^{ENOB} \cdot \min(Fs, 2 \cdot ERBW)} \quad (2.7)$$

In this equation P equals power, $ENOB$ equals the number of effective bits measured for a DC input signal, Fs equals the sampling rate, and $ERBW$ is the effective resolution bandwidth. The $ENOB$ is calculated as

$$ENOB = \frac{SINAD - 1.76}{6.02} \quad (2.8)$$

where the SINAD is measured for a (near) DC input. The effective resolution bandwidth is equal to the frequency that results in a 3 dB SINAD reduction compared to the SINAD at DC. The unit of the FoM of Eq. 2.7 is Joules per conversion step. As a result, a lower value is better. Sometimes the inverse of Eq. 2.7 is used such that a higher FoM number represents a better result.

Although the FoM of Eq. 2.7 is widely used, it cannot be used to make fair comparisons between low resolution and high resolution AD converters. When the resolution of an ADC is increased, a point is reached where thermal noise is limiting the SNR. In order to reduce the impact of the noise by 3 dB, capacitances need to be doubled. To increase the number of effective bits by one, a 6 dB reduction of the noise is required, which means a factor four increase in capacitance. Since power scales linearly with the amount of capacitance to charge, the power will also increase with a factor four. Thus, the FoM will become at least a factor 2 worse when the *ENOB* is increased by one. To enable the comparison of different resolution AD converters, an alternative version of the FoM is therefore sometimes used:

$$FoM = \frac{P}{2^{2 \cdot ENOB} \cdot \min(F_s, 2 \cdot ERBW)} \quad (2.9)$$

The equation is identical to Eq. 2.7, except that the denominator becomes four times larger instead of two times when the *ENOB* is increased by one.

Whereas comparison of AD converters by means of a single FoM is common practice, for DA converters it is not a standard approach. One of the main reasons why for DACs the single FoM approach is problematic is the time continuous output signal. When the DAC output signal is switching, i.e. making a transition between two levels, it can follow any trajectory before the signal settles to the correct value. Deviations from the ideal switching trajectory will add noise and distortion to the output. Depending on the type of application, these glitches could be problematic but not necessarily. In some applications only the DC transfer is important whereas in other applications the signal quality over a large bandwidth is important. Sometimes a signal overshoot at a transition is allowed, sometimes a smooth settling curve without overshoot is required. However, avoiding time domain glitches will typically cost power, and therefore the power efficiency of a converter can vary greatly depending on the time domain behavior.

Another reason why the single FoM approach is difficult to apply to DACs, is that part of the power consumption of a DAC is useful, and not overhead as in the case of an ADC. The output signal of a DAC is not only an information signal, but at the same time a power signal. Typically the DAC output drives a 50 Ω or 75 Ω load. If a larger output swing is required from the DAC, more power will have to be spent in the generation of this signal. A higher power consumption is thus not necessary equal to less performance, but could also indicate more performance.

In conclusion, for comparing DAC performance sometimes the FoM of Eq. 2.7 is used, but no actual de facto standard exists. However, since part of the power consumption is, by definition, required to drive the load, straightforward application of Eq. 2.7 can lead to incorrect conclusions. Other FoM measures used for DAC characterization include the SFDR, THD, and SNR, but also the static differential non-linearity (DNL) and the integrated non-linearity (INL), as well as time domain glitch energy measures.

2.4.2 *SDM Specific Functional Performance*

The SDM specific functional performance indicators relate to the stability of the converter, limit cycle and idle tone behavior, noise modulation, and transient performance.

2.4.2.1 **Stability**

Higher order Sigma-Delta Modulators are conditionally stable. As a result, only signals below a certain maximum input level can be converted without causing the modulator to become unstable. This level for which the modulator becomes unstable is a function of the loop-filter order and loop-filter cutoff frequency [49]. If the loop filter is fixed, the maximum input amplitude can be determined by means of simulations.


The procedure consists of repeatedly applying a signal with a constant amplitude to the converter. The converter is run until instability is detected or until a maximum amount of time has passed. If no instability is detected within the predetermined amount of time, it is concluded that the converter is stable for the applied signal level and the signal amplitude can be increased. If instability was detected the maximum level that can be applied has been found. Instead of trying to detect instability while the converter is running, it is also possible to always run the converter for the maximum amount of time, and afterwards determine if the converter is still stable.

With the second approach it is easier to quantify the result and this is therefore the approach taken in this work. Instability can be detected by testing the output bitstream for long sequences of 1's or 0's (hundreds of equal bits), or by testing if the modulators internal integrator values are above a certain, empirically determined, threshold. The easiest procedure is to test the output bitstream. Alternatively, the SNR and the frequency of the output signal can be measured and (near) instability can be detected by testing if the obtained values differ strongly from the expected values. This is the approach taken in this work.

Instead of measuring the maximum input signal that can be handled by the modulator, it is possible to measure how aggressive the loop filter can be made before instability occurs for a given input signal. A practical test signal is a sine wave with the maximum desired amplitude. The same procedure for detecting instability as explained above can be used, i.e. the modulator is ran for a fixed amount of time and afterwards it is determined if instability was reached. Aggressiveness of a loop filter can be increased by increasing the order of the filter or by increasing the corner frequency of the filter. Changing the filter order has a very large impact on the stability of the modulator and is therefore not practical. The loop-filter corner frequency on the other hand can be adjusted in very fine steps and is therefore more appropriate for determining stability.

In the case of a traditional SDM the stability can be determined for a given configuration, but cannot be changed or influenced in any way. For the look-ahead modulator structures in this book the situation is slightly different, and as a function of

Fig. 2.21 Illustration of the definition of a limit cycle (*left*) and an idle tone (*right*)



the available computational resources the stability will vary. It is considered beneficial to have a stable modulator to enable a large input range and high SNR.

2.4.2.2 Limit Cycles and Idle Tones

Because of the non-linear behavior of a few-bit SDM, the output signal can sometimes contain correlated frequency components that are not present in the input signal and that are not part of the normal quantization noise floor. We distinguish those components between limit cycles and idle tones. A limit cycle is a sequence of P output symbols, which repeats itself indefinitely. As a result the output spectrum contains only a finite number of frequency components. An idle tone is a discrete peak in the frequency spectrum of the output of an SDM, which is superposed on a background of shaped quantization noise. Hence, there is no unique series of P symbols which repeats itself [48]. The two situations are illustrated in Fig. 2.21.


When limit cycles are present in the output signal of an SDM, typically no signal content except DC is present at the input, although in theory also a generic repetitive input signal, e.g. a sinusoid, could be present. In practice, limit cycles only show up when the input signal is removed and a small DC offset remains. Depending on the DC level, which determines the frequency content of the limit cycle, the limit cycle can contain in-band components and cause problems or only contain harmless high frequency components.

Idle tones on the other hand typically occur when an input signal is present at the input of an SDM. Harmless high frequency idle tones are often present in the output spectrum of an SDM, but depending on the input signal the frequency of an idle tone can also be in-band and cause significant degradation of the output signal quality. Higher order modulators typically show less idle tones than low order modulators. By dithering a modulator mildly the power in idle tones can be reduced, but to fully avoid all possible idle tones a very significant amount of dither is required, penalizing the stable input range of the modulator severely. Therefore, a modulator that does not introduce idle tones or limit cycles is preferred.

2.4.2.3 Noise Modulation

Noise modulation is the effect where the amount of quantization noise in the output varies as a function of the input signal power. This effect is fundamentally present for 1-bit converters. The total output power of a 1-bit SDM is constant and equals

Fig. 2.22 SNR as a function of the input level. Ideal behavior, expected behavior, and actual measured behavior are indicated



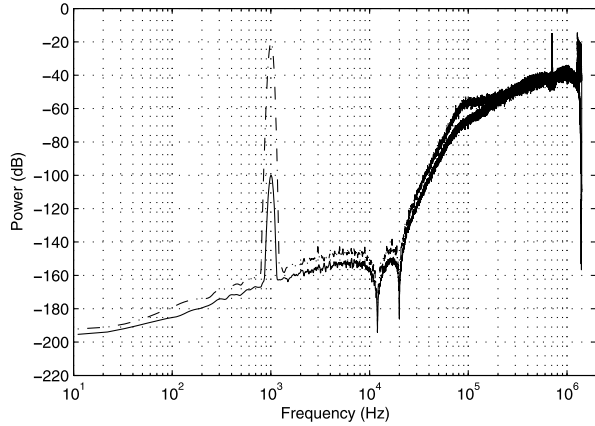
1.0, independent of the input signal power. Since the output signal power equals the input signal power, a varying amount of the output power is available for quantization noise, and it is clear that noise modulation is required to have a functional system. Therefore noise modulation in its basic form is not an issue according to the author. The problem however is located in the fact that the amount of baseband quantization noise may vary with the input signal power.

In the case where the converter is used in an audio application, and the background noise (the quantization noise) grows stronger and weaker with changes in the music level, the effect has proven to be audible in critical listening situations. According to [13, 36, 52] the variation in background quantization noise should be less than 1 dB in order to be inaudible. For high quality audio applications the objective is to have a constant background noise which results in predictable signal quality. Therefore noise modulation should be minimized or avoided if possible.

In the special case where the converter is used in a test or measurement setup and the only concern is to maximize the SNR for every input (AD) or output level (DD for DA), noise modulation can be used to an advantage. Since the total output power for a 1-bit converter is constant, the noise power increases when the signal power reduces. If the increase in noise power would be evenly distributed over all frequencies, the SNR in the baseband would decrease relatively more when the input power is reduced. In practice however, the amount of baseband quantization noise reduces when the input signal becomes smaller, and therefore the SNR will be higher than expected. This SNR behavior as a function of the input amplitude is depicted in Fig. 2.22. The SDM used in this experiment was a fifth order with two resonators. The measured, ideal and expected SNR curves are aligned to read the same value for a -25 dB input signal. The difference between the ideal and expected curve is approximately 0.7 dB for a -5 dB input, and negligible for inputs below -10 dB.

The explanation of this phenomenon can be found by studying the high frequency noise spectrum. When low amplitude inputs are applied, the SDM will start to generate high frequency idle tones, which take most of the noise power. When the input amplitude is increased, the idle tones cannot exist anymore and the low frequency noise-floor will increase [48]. In Fig. 2.23 the output spectrum for a -100 dB input

Fig. 2.23 SDM output spectrum for a -100 dB and a -20 dB input. The baseband noise-floor for the -100 dB input is lower than that of the -20 dB input because more noise power is present at high frequency in the form of idle tones



is shown in combination with the output spectrum for a -20 dB input. The baseband noise-floor of the -100 dB signal is significantly lower. Around $F_s/4$ strong tones are visible for the -100 dB input, whereas the -20 dB input has a small number of tones around $F_s/2$.

Testing for noise modulation is typically done by applying several DC levels as input signal, and for each DC level low-pass filtering the output and calculating the second order moment M_2 of the error. The advantage of this method is that it will measure exactly how much noise is generated for each input level. As an alternative it is possible to sweep the input level of a sine wave and to calculate the SINAD for each level. This method results in less precise results, since the sine wave passes through a range of intermediate levels, causing an average noise level. Although the amount of information obtained by performing a DC sweep is larger, the result from the AC sweep is more representative for specifying audio encoding quality. Both methods are used in this book.

2.4.2.4 Transient Performance

Because an SDM is an oversampled system that relies on noise shaping and feedback to realize amplitude resolution, it is not under all conditions able to encode the input signal with an equally high precision. For example, when a modulator is close to instability it can have difficulty to accurately follow transients in the input signal. When this happens, temporarily relatively large encoding errors are introduced until the modulator has recovered. Since the occurrence of this effect depends on the state of the system, it is difficult to detect or measure the impact using steady-state signals. By performing an analysis on dynamically changing signals, using a transient signal analysis method, it is possible to detect such encoding errors if they are not masked by other encoding imperfections. However, the measurement of performance in the time domain is not common practice. Therefore, a transient signal analysis method is introduced in Chap. 3.

2.4.3 *SDM Specific Implementation Costs*

For an SDM the implementation costs can be specified in more detail than for a generic converter, i.e. they can be associated with the specific SDM building blocks.

In the case of an SDM ADC, the converter consists of a loop filter, a quantizer, and a feed-back DAC. All blocks generate noise and contribute to the final SNR. The loop filter determines the order and amount of quantization noise shaping. A higher order filter will require more components, more power, and more silicon area. Since the first integrator stage of the filter typically consumes the most power, the impact of increasing the filter order is often limited. Next, there is the quantizer. If more quantization levels are required, the complexity of this block will increase. Most commercial designs do not have more than 5 quantization bits, because the design efficiency will go down strongly for more bits. This efficiency reduction is caused by the increase of detection levels and at the same time the more stringent requirement on accuracy and noise performance of the thresholds. In most designs the number of quantization levels of the DAC equals that of the ADC. Thus, an increase in complexity of the quantizer also enforces an increase in the complexity of the DAC. Another factor that influences the cost of the building blocks is the oversampling ratio (OSR) of the converter. An increase of the OSR will require the quantizer to make its decisions faster and the DAC to produce its output faster, i.e. they run at a higher clock frequency. In most situations this will result in an increase in power consumption of those blocks. Furthermore, the output data rate of the converter also scales with the OSR, and the complexity of the subsequent digital decimation stages will also be affected. However, a higher OSR is beneficial for the noise-shaping efficiency, and might allow for a lower order filter or less quantization levels without reducing the SNR of the converter.

In the case of a digital-to-digital SDM the implementation costs are different. All functions, i.e. the loop filter and quantizer, are realized by digital logic functions. The loop filter is by far the most difficult block and requires significant hardware resources. The filter consists of integrators, which are digitally realized by adders with feed-back around a delay element, and filter coefficients. The coefficients are realized by digital multipliers or combinations of shift and add. Because typically a large dynamic range is required from the modulator, a very wide data-path is required. On the other hand, the quantizer function often involves not more than a few comparators. In the case of a 1-bit SDM, it can even be implemented without any comparison operations, and selection of the sign bit is all that is required. Also in the case of a DD converter it is the OSR of the converter that determines the clock frequency at which all the operations are performed. For most digital realizations the power consumption scales with the clock frequency. Initially this scaling is linear, but if the frequency is increased towards the limit of the technology, the scaling is more than linear. Because an SDM is a feed-back structure, pipelining of operations is not possible and all the results should be ready within a single clock cycle. The computation of the coefficient multiplication is the most challenging operation, but by selecting appropriate coefficients the computations can often be simplified.

Next to the already mentioned complexity of the loop filter, the quantizer, and possibly the feed-back DAC, another source of complexity exists for the Sigma-Delta Modulators discussed in this book, namely the complexity resulting from the addition of look-ahead. Without going into the details of the look-ahead concept (see Chap. 5 and further), the complexity can be summarized as follows. In order to realize a look-ahead modulator, a multitude of loop filters is required. For each loop filter an alternative quantizer is present. Finally, there is a control structure that takes care of the selection of the output symbol. Because of the multiple loop filters and quantizers, the power consumption of a look-ahead modulator will be a multiple of that of a normal modulator. On top of the increased hardware complexity, there is also an increased algorithmic complexity associated with the look-ahead concept.

2.4.4 Figure-of-Merit of an SDM

The efficiency of an SDM is typically compared using an FoM, just as is done for any other data converter. Depending on the type of SDM, i.e. an ADC, a DAC, or a DD converter, the FoM used is different.

In the case of an ADC SDM typically only the power consumption of the analog part is measured when the FoM is calculated, i.e. the digital decimation filter is often not considered in the efficiency. It is the view of the author that this is not correct, since the reconstruction filter is essential for the operation of the modulator, and can consume a considerable amount of power. Nevertheless, in practice the FoM is mostly calculated with a slightly modified version of Eq. 2.7 (or Eq. 2.9):

$$\frac{P}{2^{ENOB} \cdot \min(BW, ERBW)} \quad (2.10)$$

In this equation BW equals the bandwidth that is used in the SNR calculation. In this equation the conversion bandwidth is thus limited to the smallest of the $ERBW$ and the signal conversion bandwidth. Besides this change, the FoM equation is identical to the generic one and no SDM specific features are included.

For an SDM based DAC the calculation of an FoM is even more dubious than for a generic DAC. Without the digital-to-digital converter that drives the DA stage, the DAC cannot work. Therefore, only by including the power of the digital SDM a sensible FoM can be calculated. The problem of selecting an appropriate FoM is now similar to the situation of a generic DAC, and Eq. 2.10 is typically used.

In the case of a stand-alone DD converter, FoM calculations like the one of Eq. 2.10 are typically not used. Most stand-alone converters are software based instead of dedicated hardware solutions, and therefore the power measure is not practical. A convenient metric in this case is the amount of operations per second required for the implementation to run real-time. Alternatively, the absolute amount of time can be measured that is required to process a fixed amount of signal. If the same test conditions are used repeatedly, i.e. same test signal and same computer platform, results can be compared and a valid FoM measure can be derived.

Although the above mentioned methods can be conveniently used to measure the computational efficiency of a DD design, no signal conversion performance is included in this FoM. This is not a problem if the designs under comparison are designed to deliver equal performance, i.e. have the same loop filter and have the same OSR. If the resulting signal conversion performance is different for the converters under comparison, only measuring the computational efficiency is not good enough. However, an FoM like Eq. 2.10 where P is replaced by computational load cannot be used since there is not a direct relation between the signal conversion performance and the computational load, i.e. a change of the loop-filter transfer will affect the SNR but not necessarily the amount of computations.

Because of the issues in defining a simple FoM that includes all the relevant measures, stand-alone DD Sigma-Delta Modulators will have to be compared using several measures, similar to the situation of comparing generic DACs. For designs that realize an identical SNR, the relevant signal quality measures are the SFDR and the THD. Next to these generic measures, designs can be compared on the relevant SDM specific measures, mainly the stability of the converter, the transient behavior, and possibly the amount of noise modulation. The realized signal quality can then be compared with the computational load required for this quality.

Deciding on what design is the best on the basis of the different metrics, however, is not straightforward if similar performance levels are reached. In this case the design with the lowest computational load is selected as the better one. For practical reasons, in this book the computational load of the different DD converters is measured by recording the time that is required to process a specific test signal. These measurements are performed on the same computer platform under the same conditions, such that the results can be used to classify performance levels. Note that the impact of the computer architecture is not considered in this performance evaluation, and that the results cannot be used for generic benchmarking purposes against literature.



<http://www.springer.com/978-94-007-1386-4>

Look-Ahead Based Sigma-Delta Modulation

Janssen, E.; van Roermund, A.

2011, XII, 248 p., Hardcover

ISBN: 978-94-007-1386-4