ACSP · Analog Circuits And Signal Processing

Erwin Janssen
Arthur van Roermund

Look-Ahead Based Sigma-Delta Modulation

Springer

Look-Ahead Based Sigma-Delta Modulation

ANALOG CIRCUITS AND SIGNAL PROCESSING

Series Editors:
Mohammed Ismail. ***The Ohio State University***
Mohamad Sawan. ***École Polytechnique de Montréal***

For other titles published in this series, go to
www.springer.com/series/7381

Erwin Janssen • Arthur van Roermund

Look-Ahead Based Sigma-Delta Modulation

Springer

Erwin Janssen
NXP Semiconductors, Central R&D
Mixed Signal Circuits & Systems
High Tech Campus, Building 32
Eindhoven 5656 AE
The Netherlands
erwin.e.janssen@nxp.com

Arthur van Roermund
Electrical Engineering
Mixed-signal Microelectronics Group
Technical University Eindhoven
Eindhoven 5600 MB
The Netherlands
a.h.m.v.roermund@tue.nl

Series Editors:
Mohammed Ismail
205 Dreese Laboratory
Department of Electrical Engineering
The Ohio State University
2015 Neil Avenue
Columbus, OH 43210
USA

Mohamad Sawan
Electrical Engineering Department
École Polytechnique de Montréal
Montréal, QC
Canada

ISBN 978-94-007-1386-4 e-ISBN 978-94-007-1387-1
DOI 10.1007/978-94-007-1387-1
Springer Dordrecht Heidelberg London New York

Library of Congress Control Number: 2011926230

Cover design: VTeX UAB, Lithuania

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Contents

List of Symbols and Abbreviations

t_0	current time step
AC	alternating current
AD	analog to digital
ADC	analog to digital converter
C	accumulated cost value or cost function
c	cost value
DA	digital to analog
DAC	digital to analog converter
dBFS	decibels full scale
DC	direct current, 0 Hz
DD	digital to digital
DDC	digital to digital converter
DSD	Direct Stream Digital
DSM	Delta-Sigma Modulator
DST	Direct Stream Transfer
ENOB	effective number of bits
ERBW	effective resolution bandwidth
ETSDM	Efficient Trellis SDM
FB	feed-back
FF	feed-forward
FoM	figure of merit
Fs	sampling rate
HD	harmonic distortion
HMM	hidden Markov model
L	the number of bits latency, or the trace-back depth in bits, for all the different look-ahead algorithms
LA	look-ahead
LASDM	look-ahead SDM
M	the number of parallel paths in the Pruned Tree sigma-delta modulation algorithm (for SA-CD)

N	the number of bits over which uniqueness of the parallel solutions is determined in the (Efficient) Trellis sigma-delta modulation algorithm. Also determines the number of parallel solutions (2^N) for the Trellis algorithm
NS	noise shaping
NTF	noise transfer function
OSR	oversampling ratio
PCM	pulse-code modulated
PDF	probability density function
PDM	pulse density modulation
PTSDM	Pruned Tree SDM
PWM	pulse width modulation
SA-CD	Super Audio CD
SD	Sigma-Delta
SDM	Sigma-Delta Modulator
SFDR	spurious free dynamic range
SINAD	signal to noise and distortion ratio
SNDR	SINAD
SNR	signal to noise ratio
STF	signal transfer function
TD	time domain
THD	total harmonic distortion
TSDM	Trellis SDM

Chapter 1
Introduction

In March 1999 the Super Audio Compact Disc (Super Audio CD, SA-CD), the successor of the normal audio CD, was presented to the world. This new audio carrier, conceived by Philips and Sony, makes use of a radically new way to store and reproduce audio signals. Instead of working with the traditional 44.1 kHz sampling rate and 16-bit pulse-code modulated (PCM) signals, a 2.8 MHz 1-bit format is used to store the audio signal. The new format is marketed to deliver a signal-to-noise ratio (SNR) of 120 dB and a signal bandwidth of 90 kHz, as opposed to an SNR of 96 dB and a bandwidth of 20 kHz for the normal audio CD. The decision for this alternate encoding format was made years earlier, when 1-bit Analog-to-Digital (AD) audio Sigma-Delta (SD) converters were still delivering the highest signal conversion quality. In fact, virtually all of the high quality AD and Digital-to-Analog (DA) converters that were used at that time for the generation and reproduction of CD quality PCM audio were based on 1-bit converters. It was reasoned that a higher audio quality could be obtained by removing the decimation and interpolation filters that performed the conversion from 1-bit to PCM and vice versa, and by storing the 1-bit signal from the Sigma-Delta Modulator (SDM) directly on the disc.

Although the idea of storing the 1-bit SDM output signal directly on the disc sounds very reasonable, in practice things work differently, and the original recorded signal is never stored directly on a disc. What typically happens is that a number of recordings of the same performance are made, and that at a later stage in the studio those recordings are edited and processed, e.g. removal of coughs from an audience or the equalization of the audio levels, until the desired sound quality is obtained. This process of editing and processing can only be performed on multi-bit (PCM) signals, and only once all this work is done the 1-bit signal that will be stored on the SA-CD disc will be generated. Thus, if it is assumed that all the digital processing on the audio signal is without any loss of the signal quality, the final signal quality of the 1-bit signal that is stored on the disc is determined by the initial analog-to-digital conversion and the final Digital-to-Digital (DD) conversion.

Nowadays, the highest quality analog-to-digital conversion for audio applications is obtained with a multi-bit SDM. Such a converter can deliver a very high SNR and very low distortion levels. From the output of the SDM a PCM signal is generated,

E. Janssen, A. van Roermund, *Look-Ahead Based Sigma-Delta Modulation*,
Analog Circuits and Signal Processing,
DOI 10.1007/978-94-007-1387-1_1, © Springer Science+Business Media B.V. 2011

but now with a higher resolution and much higher sampling rate than what is used for CD. After all the processing on the multi-bit signal is performed, the final 1-bit signal is generated. Traditionally, this is done with a digital 1-bit SDM. However, with a normal SDM it is not trivial to generate a 1-bit signal with the desired ultra-high quality under all signal conditions. For example, for extremely high signal levels a 1-bit SDM can generate significant distortion, especially if the modulator is designed to deliver a very high SNR for normal signal levels. Besides this potential signal quality issue there is a much bigger issue that, with traditional sigma-delta modulation approaches, cannot be solved without jeopardizing the signal quality: the risk of not realizing a long enough playback duration.

The SA-CD standard supports, in addition to a normal stereo recording, also the possibility to store a multi-channel version of the same recording. In order to fit all the data on the 4.7 gigabyte disc and obtain a playback duration of at least 74 minutes, the standard playback duration of the normal audio CD, lossless data compression is applied to the 1-bit audio signal. Only if the compression gain, the ratio that indicates the amount of data size reduction, is high enough it will be possible to obtain the required 74 minutes of playback time. Since the data compression algorithm is lossless, the compression gain depends on the redundancy in the 1-bit encoded audio signal, and this can only be influenced with the SDM design. However, the only solution to increase the redundancy is to reduce the signal conversion quality of the SDM, and since SA-CD is about delivering high audio quality this is not an acceptable solution.

In order to come to an efficient solution to the above mentioned problems, the possibilities and opportunities of look-ahead based sigma-delta modulation are discussed in this book. This exploration is performed with a focus on generic 1-bit look-ahead sigma-delta modulation, applicable to any sampling rate and loop filter type. No use is made of SA-CD specific nomenclature, except for Chap. 10 where a minimal amount of usage cannot be avoided.

In Chap. 2, a basic introduction to sigma-delta modulation and the performance evaluation of Sigma-Delta Modulators is given. Readers familiar with traditional sigma-delta modulation for AD and DD conversion and the possible artifacts resulting from 1-bit sigma-delta modulation can skip this chapter and immediately continue with Chap. 3.

Traditionally, signal conversion quality is characterized with steady-state signals. In the case of a linear data converter this procedure will also give the performance for non-steady-state signals. However, since a 1-bit SDM is a non-linear data converter, it is not guaranteed that the steady-state performance is representative for non-steady-state signals. In Chap. 3 this potential discrepancy is investigated.

In Chap. 4, a generic model of a noise-shaping quantizer is derived. This model is subsequently used in Chap. 5 to come to a noise-shaping quantizer model for a look-ahead converter. Next, the main look-ahead principles are introduced, accompanied with an analysis of the benefits and disadvantages. The basic full look-ahead algorithm is presented, and an analysis is made of the possibilities for realizing a look-ahead enabled AD converter. Although this idea is rejected, it is clear that large benefits can be expected from look-ahead based DD conversion, but only if an approach with a reduced computational load can be realized.

The possibilities for reducing the computational load of the full look-ahead algorithm for DD conversion are investigated in Chap. 6. Since the obtainable reduction is rather limited, an alternative approach, i.e. pruning of the solution space, is investigated. It is concluded that, with a proper pruning algorithm, it should be possible to realize solutions that result in large computational savings and that have a limited impact on the obtainable signal conversion performance. Therefore, the next chapters focus on pruned look-ahead algorithms.

In Chap. 7, an analysis is made of the Trellis sigma-delta modulation algorithm by Kato. An improvement of the signal conversion quality, compared to a normal SDM, is realized but at a very large computational cost.

Further analysis of the Trellis sigma-delta modulation algorithm in Chap. 8 reveals that only a fraction of all the parallel solutions contributes to the final output. The Efficient Trellis sigma-delta modulation algorithm makes use of this observation and prunes the solution space further, thereby enabling a larger pruned look-ahead depth that results in an improvement of the signal conversion quality, as well as a reduction in the computational load.

In Chap. 9, the Pruned Tree sigma-delta modulation algorithm, that is an improvement over the Efficient Trellis sigma-delta modulation algorithm, is discussed. The pruning criteria that is applied in the Efficient Trellis sigma-delta modulation algorithm is effective for reducing the number of parallel solutions, but also adds a significant computational overhead to the algorithm. By changing the initial conditions of the look-ahead modulator the pruning criteria can be relaxed, which results in a computationally more efficient solution that, typically, delivers performance that is on par with that of the Efficient Trellis sigma-delta modulation algorithm, but that is sometimes even better.

In the Pruned Tree sigma-delta modulation algorithm for SA-CD, described in Chap. 10, a cost function is added to the original Pruned Tree sigma-delta modulation algorithm that reflects the predictability of the output bitstream. This addition results in a dual optimization that takes both the signal quality into account and improves the lossless data compression gain of the output signal.

In Chap. 11, a comparison is made between the various look-ahead techniques that are detailed in the previous chapters. This comparison includes an analysis of the algorithmic differences, and a comparison of the functional performance.

In the previous chapters it was found that there appears to be a limit on the SNR that can be achieved with a fifth order 1-bit SDM. In Chap. 12 this phenomenon is analyzed in detail and new results on the limits of 1-bit noise shaping are presented.

Finally, in Chap. 13 the general conclusions on the work described in this book are presented.

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

E. Janssen, A. van Roermund, *Look-Ahead Based Sigma-Delta Modulation*,
Analog Circuits and Signal Processing,
DOI 10.1007/978-94-007-1387-1_2,

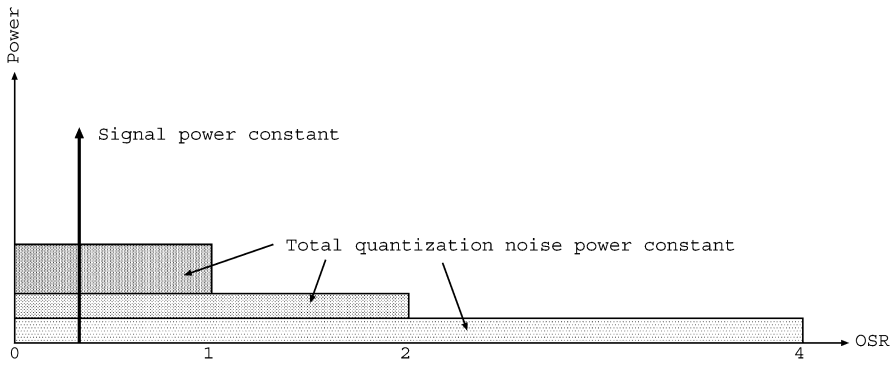

Fig. 2.1 Oversampling does not affect the signal power or total quantization noise power but reduces the noise spectral density

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

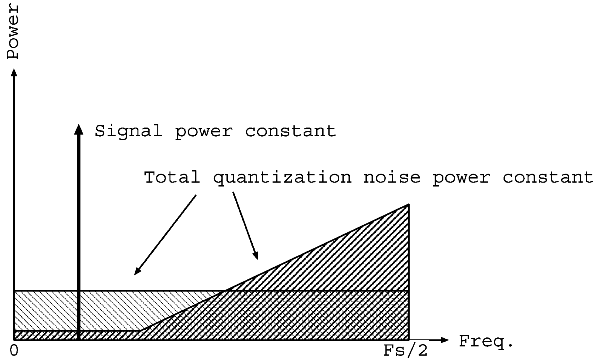

Fig. 2.2 Low frequency noise is pushed to high frequencies by noise shaping

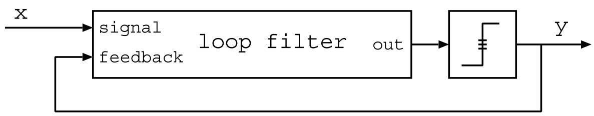

Fig. 2.3 Generic model of the Sigma-Delta noise-shaping loop, consisting of 2-input loop filter and quantizer

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.

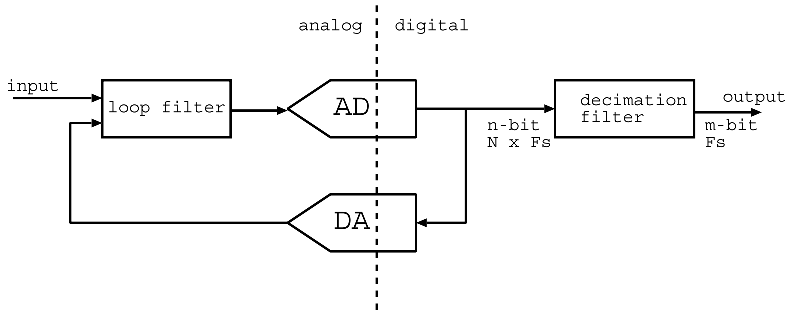

Fig. 2.4 Main building blocks of a Sigma-Delta analog-to-digital converter

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.

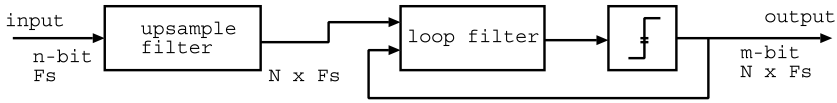

Fig. 2.5 Main building blocks of a Sigma-Delta digital-to-digital converter

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 Fs to $N \times Fs$ 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 $N \times Fs$ 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.

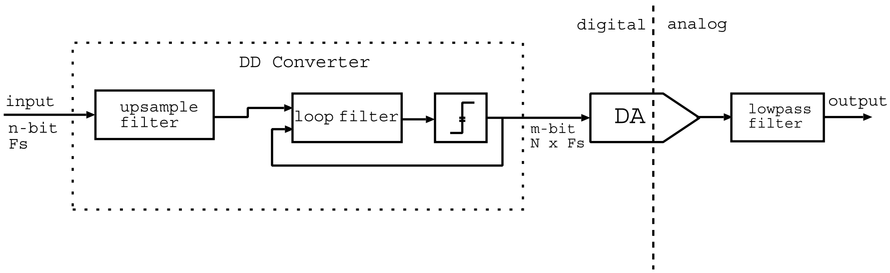

Fig. 2.6 Main building blocks of a Sigma-Delta digital-to-analog converter

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

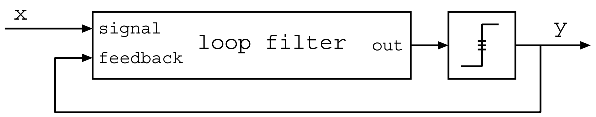

2.2 Sigma-Delta Structures

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.

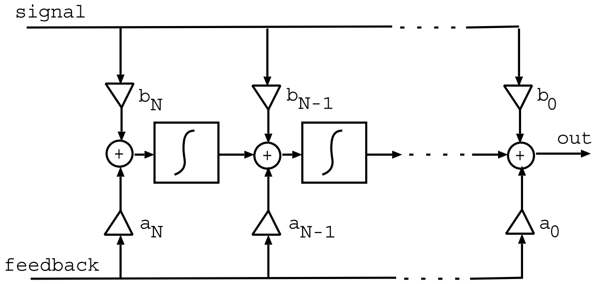

Fig. 2.8 Internal structure of practical 2-input loop filter, consisting of integrators, subtraction points, feed-forward coefficients b_i and feed-back coefficients a_i

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

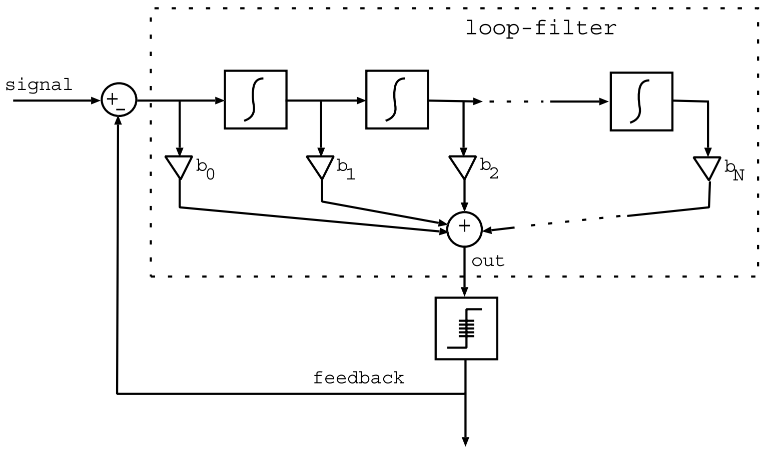

Fig. 2.9 SDM with feed-forward loop filter. The subtraction point of signal and feed-back has been shifted outside the loop filter

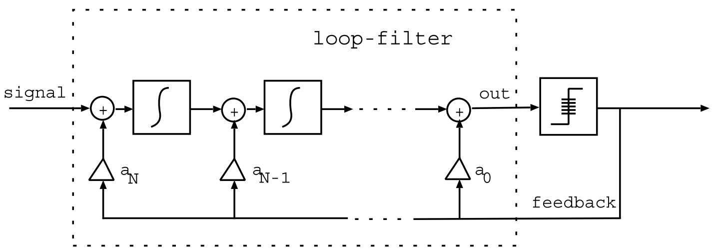

Fig. 2.10 SDM with feed-back loop filter

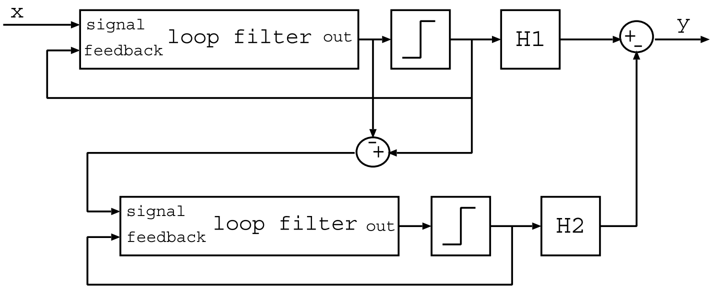

Fig. 2.11 Second order MASH SDM

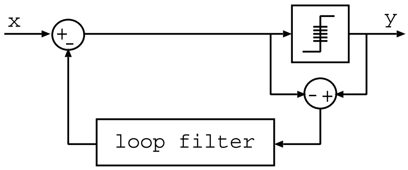

Fig. 2.12 Noise shaper structure

and only the shaped noise of the second modulator remains. In this fashion an nth 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.

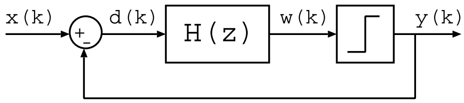

Fig. 2.13 Generic model of a digital SDM in feed-forward configuration

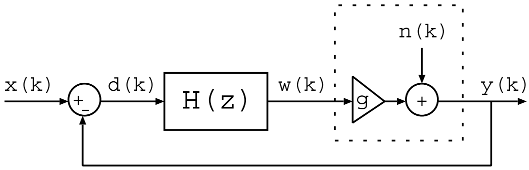

Fig. 2.14 Linear model of a digital SDM in feed-forward configuration

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:

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

$$y(k) \cdot \left[1 + H(z)\right] = H(z) \cdot x(k) + e(k) \tag{2.2}$$

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

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

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

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

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

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:

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

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.

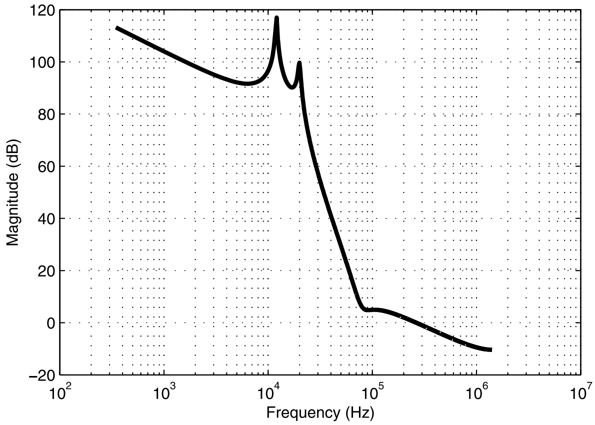

Fig. 2.15 Transfer of a typical fifth order loop filter designed according to a Butterworth specification with 100 kHz corner frequency and additional resonator sections at 12 and 20 kHz. The sampling rate is 2.8 MHz

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 $Fs/2$ 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