Unleash the Performance Benefits of Sigma-Delta ADCs into Your SoC

IP supports cellular communications, sensors and measurement markets

December, 2011

Author

Carlos Azeredo-Leme
Analog Design,
Senior Staff,
Synopsys, Inc.

Overview

Sigma-delta analog-to-digital converters (ADCs) deliver high resolution with low silicon area and power consumption, taking advantage of high-speed operation and digital signal processing common in modern digital system-on-chip (SoC), making them ideal for implementations in deep sub-micron processes. However designers are sometimes reluctant to use sigma-delta ADCs because of their unconventional architecture and apparently chaotic internal operation.

This paper will give SoC designers a clear understanding of the workings inside sigma-delta ADCs and explain when they are a better alternative than other ADC architectures used in advanced, deep sub-micron SoCs. The paper will address the following with regards to sigma-delta ADCs:

▸ Fundamentals and operating modes
▸ Key features and cost considerations
▸ How to select the right ADC architecture to meet the performance requirement of the SoC
▸ How integrating sigma-delta ADCs can offer significant advantages in a variety of applications

By understanding the benefits of the sigma-delta ADCs, SoC designers are better equipped to make design decisions that lead to compelling products that deliver significant performance, area and power gains.

The Fundamentals of Sigma-Delta ADCs

ADCs operate by comparing the input signal with internal references representative of the possible output digital codes. As the resolution increases, the number and accuracy of these references increases exponentially. For example, for 16-bit resolution, there are 65536 possible output codes each with an accuracy of 0.0015%. That is very difficult to achieve in modern SoCs due to the strong constraints in silicon area. Sigma-delta ADCs operate differently. They exploit the high-speed and large digital computation capability of today's deep submicron process technologies to extract accuracy.

A typical block diagram of a sigma-delta ADC is shown in Figure 1. It is composed of two main blocks, namely a sigma-delta modulator and a digital filter.
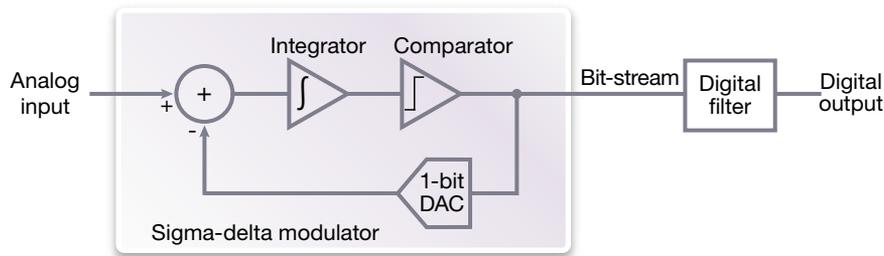
Figure 1: Block diagram of a sigma-delta ADC, consisting of a sigma-delta modulator and a digital decimating filter

The sigma-delta modulator generates a controlled pseudo-chaotic behavior by putting together the following components: memory (the integrator), non-linearity (the comparator) and feedback (through the 1-bit digital-to-analog-converter (DAC)), all running at high-speeds above the input signal bandwidth. The result is a bit-stream output that jumps at high speed between logic 0 (representing the maximum negative input) and logic 1 (representing the maximum positive input). This random sequence of 0's and 1's includes all the information about the input signal plus the internally generated high-speed noise.

The resulting bit-stream for a sinusoidal analog input is shown in Figure 2. It is apparent that the density of 1's is higher during the positive half-cycle and lower during the negative half-cycle. Consequently, the density of 1's and 0's is balanced during the zero crossings, following the input sinusoidal waveform.

The function of the digital filter block is to extract the signal from the added high-frequency noise by low-pass filtering and decimation. The output is a typical N-bit word at the desired conversion rate of a typical ADC.
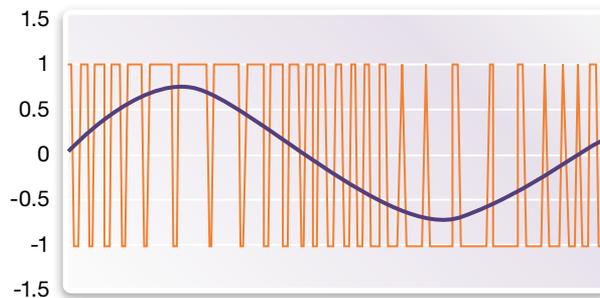
Figure 2: Resulting bit-stream of the sigma-delta modulator for a sinusoidal analog input

Sigma-Delta Modulator Operation

The process of digitalization in ADCs introduces quantization errors due to the finite output word-length, as illustrated in Figure 3. These errors are bounded by the height of the quantization step (Δ). Assuming an N-bit ADC with a normalized analog input range of 1 V, the quantization step is denoted by $\Delta = 1/(2^N-1)$ because there are 2^N codes covering the input range from -0.5 V to +0.5 V. It can be noted that such quantization errors will limit the ADC maximum Signal-to-Noise Ratio (SNR) defined as the ratio between the full-scale sinewave power and the total quantization power according to:

$$SNR=6.02N+1.76 \ [dB] \qquad \text{[1]}$$

This expression is a direct relationship between resolution and maximum SNR. Therefore, the resolution of an ADC can, alternatively, be expressed in SNR using the correspondence in Table 1.
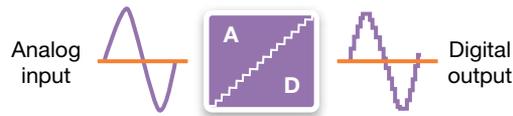
Figure 3: Illustration of the quantization effect of an ADC on an input sinewave

Resolution	SNR	Δ
1-bit	8 dB	1 V
2-bit	14 dB	0.33 V
3-bit	20 dB	0.14 V
4-bit	26 dB	0.067 V
5-bit	32 dB	0.032 V
6-bit	38 dB	0.016 V
7-bit	44 dB	7.9 mV
8-bit	50 dB	3.9 mV
9-bit	56 dB	2.0 mV
10-bit	62 dB	0.98 mV
11-bit	68 dB	0.49 mV
12-bit	74 dB	0.24 mV
13-bit	80 dB	0.12 mV
14-bit	86 dB	61 uV
15-bit	92 dB	31 uV
16-bit	98 dB	15 uV

Table 1

Using this concept of SNR as an equivalent specification for the resolution of an ADC, we can investigate new ways to increase the resolution without relying on refining the quantization step (Δ):

▸ One possible way is to oversample, meaning to clock the ADC above the Nyquist rate. The Nyquist rate is the minimum conversion rate to reconstruct a band-limited signal and corresponds to twice the signal bandwidth. In fact, the quantization noise power is only a function of the quantization step (Δ) and not of the conversion frequency. Therefore, any quantization noise outside the signal band can be removed with a digital filter after the ADC, thereby improving SNR and consequently increasing resolution. This is illustrated in Figure 4. In Figure 4a the conversion rate (F_s) is at Nyquist, meaning that all the quantization noise is distributed between 0 and half the conversion rate, $F_s/2$, and is superimposed with the signal bandwidth, Bw.

▸ By oversampling to $F_s/2 >> Bw$, as shown in Figure 4b, only a fraction of the quantization noise falls in-band and the rest can be removed by a digital filter. With this technique the resolution can be increased by 1-bit for each 4X increase in the conversion rate (F_s).

a) Conversion at Nyquist

b) Conversion with oversampling

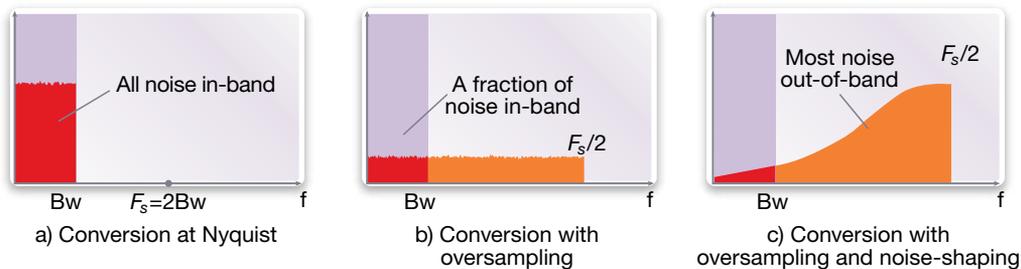c) Conversion with oversampling and noise-shaping

Figure 4(a,b,c): Effect of oversampling and noise-shaping on quantization noise

▶ Adding noise-shaping to oversampling results in the situation shown in Figure 4c. Most noise is pushed to high frequency leaving only very little noise in-band. After the digital filter removes all the out-of-band noise, the increase in SNR is substantial, leading to large increases in resolution. This important noise-shaping effect is the fundamental property of a sigma-delta modulator and it is due to the integrator inside the loop (Figure 1). The noise-shaping effect can be strengthened by adding more integrators in cascade, resulting in *higher order* sigma-delta modulators. The increase in resolution is so substantial that the internal ADC can be reduced to simply a comparator implementing a 1-bit conversion.

Figure 5 shows the resulting theoretical SNR as a function of oversampling for sigma-delta modulator of orders (L) 1 up to 5. It can be seen that a first-order (L=1) modulator provides 12-bit resolution (SNR =74 dB) for an oversampling ratio of 500X. With second-order (L=2), 16-bit resolution (SNR = 98 dB) can be achieved with 200X oversampling. A third-order (L=3) modulator would require only 50X oversampling for the same 16-bit resolution. In practice, these theoretical numbers must be relaxed a bit to take into account other noise sources (e.g. thermal) and stability considerations in the higher order modulators.
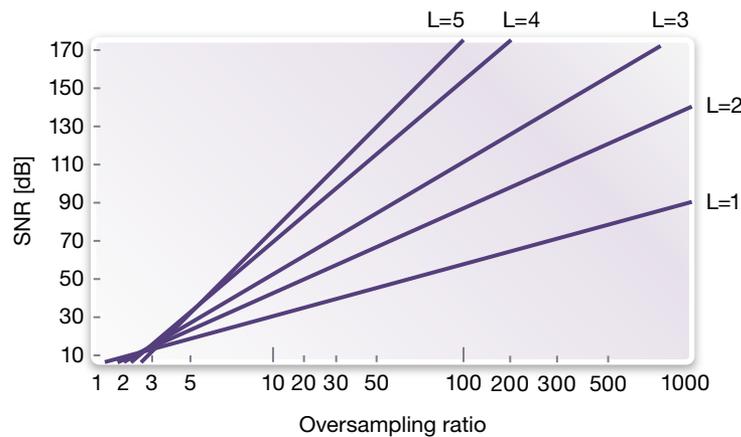

Figure 5: Theoretical SNR as a function of oversampling in sigma-delta modulators of order (L) 1 to 5

Digital Filter

The digital filter must both remove the out-of-band noise and also decimate the highly oversampled bit-stream to the output sample rate. A direct implementation would result in a very high order filter due to the large ratio between the signal bandwidth and oversample-rate. Therefore, in practice, a two-stage filter is used where the first stage just removes the higher aliasing bands and the second implements the sharp base-band response.

The first stage is typically a SINC filter. This type of filter has a very compact implementation using Cascaded Integrator-Comb[1] (CIC) structures. It conveniently introduces notches at the aliasing bands leading to an efficient decimation process. The order of the SINC must be designed to adequately remove the quantization noise from the aliasing bands so that no significant degradation of SNR results after their folding back into the base-band by the decimation process. The established practice is to design the order of the SINC filter one above the order of the sigma-delta modulator. For example, a second-order modulator requires a third-order SINC filter.

The SINC filter, however, introduces a droop in the pass-band. Therefore, it is normally not used for the complete decimation to the output data rate. The SINC filter output is typically 4X or 8X the final output rate and the remaining 4X or 8X decimation is done in the second stage as stated above.

---

[1] Eugene Hogenauer, "An Economical Class of Digital Filters for Decimation and Interpolation", IEEE Trans. on Acoustics, Speech and Signal Processing, Vol. ASSP-29, No. 2, April 1981, pp. 155-162.

Figure 6 shows a plot of the frequency response of a third-order SINC filter on a 32X oversampling system and implementing decimation by 8X. The Nyquist frequency is normalized to 1 MHz, resulting in an oversampling frequency of 32 MHz and an output rate for the SINC filter of 4 MHz. The frequency response shows notches at 4 MHz and its multiples (8 MHz, 12 MHz, 16 MHz…) that correspond to the centers of the aliasing bands of the decimation by 8X from 32 MHz to 4 MHz. Since the output rate from the SINC is still 4X the Nyquist rate, the in-band droop is small and can be easily compensated by the next filter stage.
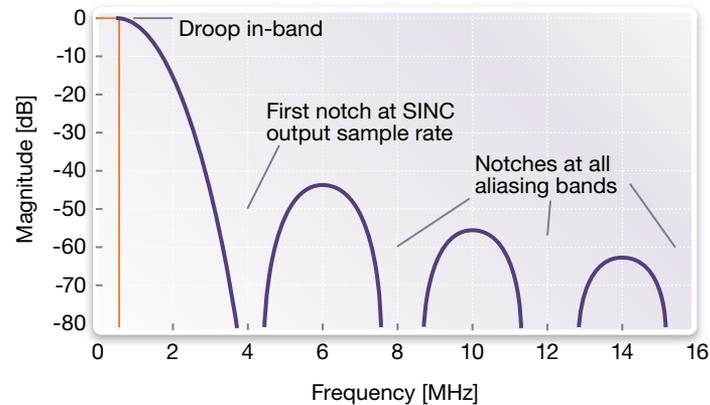
Figure 6: SINC^3 Filter plot with decimation by 8.
Nyquist rate is normalized to 1 MHz and oversampling rate is 32 MHz

The next filter stage implements the final decimation ratio to produce the output samples at the desired output rate, which can be the Nyquist rate or a higher frequency. This filter can be one of two types:

▶ Infinite Impulse Response (IIR), also known as recursive because its output depends on both past input and past output values
▶ Finite Impulse Response (FIR), also known as non-recursive because its output depends only on past input values

IIR digital filters have poles and zeros, like analog filters. The poles allow very efficient transfer functions to be realized so the order or complexity of the implementation is reduced. But the poles introduce phase response distortions. They are the most efficient choice when no linear phase response is required, such as in voice and sensor data.

FIR filters have no poles and they can be made such that all zeros are ideal notches (on the unit circle in Z-plane). These filters guarantee linear phase response. However, for achieving the same response specification, FIR filters must be of much higher order than an IIR equivalent. Even if the complexity for implementing a zero is lower than for a pole, the complexity of a FIR filter is normally higher than for the corresponding IIR, therefore, resulting in larger area and consumption. They are used in Hi-Fi audio and complex modulation communication standards to preserve phase linearity.

Key Features and Cost Considerations of Sigma-Delta ADCs

The ADCs based on sigma-delta modulators gained popularity in the late 80s when the scaling of CMOS processes started allowing compact integration of the accompanying digital decimation filters. With today's nanometer processes packing millions of gates per mm^2 this is more valid than ever. In fact, the area and consumption of the complete digital filters required is now a fraction of the modulator itself. Furthermore, their area and consumption competes favorably with the equivalent implementation of the filters in analog form.

Key features

A typical communications analog front-end for signal digitization is shown in Figure 7. The input signal is filtered and then amplified. The analog filter removes both out-of-band interferences and noises that would

otherwise be folded down in-band after the analog-to-digital conversion. The amplifier raises the signal amplitude to fit into the ADC input range. The ADC is the last block in the chain and is clocked not much above the minimum Nyquist rate to minimize power consumption.

The main issue with this architecture is the high complexity of the analog filter that leads to large implementation area. This analog filter needs to implement a low-pass frequency response with a sharp roll-off at the in-band edge, resulting in high order transfer function and large resistor/capacitor sizes. To alleviate this, the ADC is often oversampled by 2-4X allowing the analog filter to relax a bit at the expense of additional digital filters.
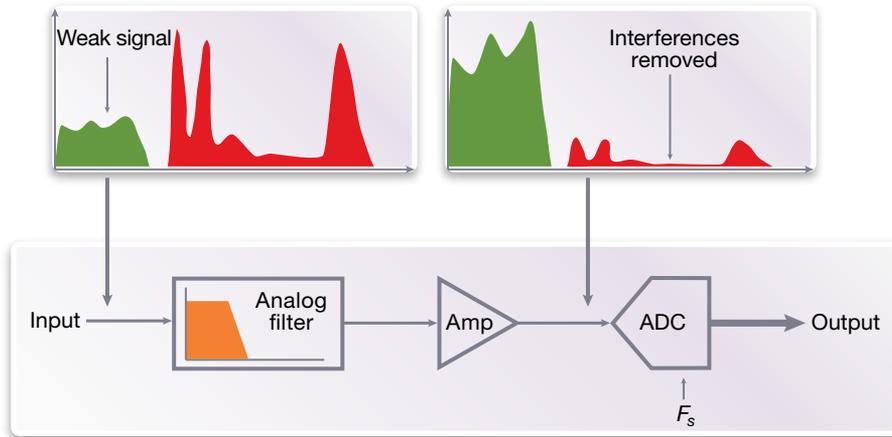

Figure 7: Classical analog front-end for signal digitization

In the case of a sigma-delta ADC, a relaxed low-pass analog filter can be used to remove the aliasing band at the high oversampling frequency (M F_s), as illustrated in Figure 8. The ADC digitizes both the signal and the out-of-band interferences together. Actually, the out-of-band energy is helpful to feed the pseudo-chaotic processes in the sigma-delta modulator avoiding periodic patterns that could result in in-band noise tones. The digital decimating filter that is part of the sigma-delta ADC for removing the quantization noise has a double function of removing the out-of-band interferers, leading to a very efficient hardware implementation.
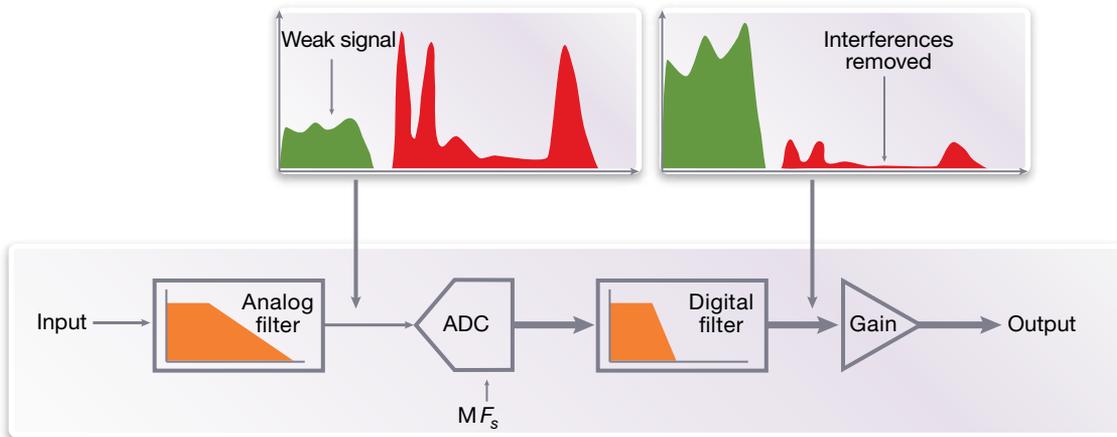

Figure 8: Sigma-delta analog front-end for signal digitization

Four important features of a sigma-delta ADC solution are:

▸ The digitized bandwidth is very wide, only limited by the high oversampling frequency. That means that the out-of-band signals can be allowed to go through the ADC and be removed afterwards by the digital decimating filter.

- The very high SNR achieved in a sigma-delta modulator leads to a very wide dynamic range. This allows digitizing even small signals in the presence of much stronger out-of-band interferences. After digital removal of the interferences by the digital decimating filter, the remaining weak signal can be scaled up to full scale by clipping the output word accordingly (digital gain).
- The elimination of the large and complex sharp roll-off analog filter required in a classical analog front-end comes at virtually no cost because a sigma-delta modulator always requires a decimating filter to remove the quantization noise and also doubles as the function of removing the out-of-band interferences. Moreover, this digital filter is free of mismatches that are typical of analog implementations, allowing mathematically-ideal transfer functions to be realized.
- Finally, sigma-delta ADC systems are robust to clock jitter and noise crosstalk due to the high oversampling frequency and minimal analog content. All the noise outside the signal bandwidth is removed by the digital decimating filter, which for a 100X oversampling ratio can be 99% of the total noise.

Cost Considerations

What are some of the costs SoC designers should consider when implementing a sigma-delta solution?

- An oversampling clock. If this is not available in the application it must be generated with a PLL.
- A digital filter to decimate the bit-stream and produce the data output.
- Higher dynamic range for the ADC: In the case, when the input signal needs amplification and that function is moved to the digital domain, the ADC must have a correspondingly lower noise floor by the same gain factor to keep the signal-to-noise ratio equivalent to the Nyquist solution.
- Loss of a direct correspondence between an input sample and the output. Indeed, the digital decimating filter produces output digital samples that are, each, a function of multiple input samples according to its impulse response. As an example, Ping-Pong architectures, where the same ADC is shared between multiple inputs such as in I&Q front-ends, are not possible with sigma-delta ADCs. Instead, one sigma-delta ADC must be used for each input.

The good news is that in many situations these costs do not apply or are part of the system implementation. In fact,

- The clock frequencies for signal processing in a modern SoC are often much higher than the Nyquist rate. It is often quite straightforward to find an appropriate oversampling clock with no need for an additional PLL.
- The digital decimating filter is often re-used to replace the analog filters that reject the out-of-band interferers and noise.
- The applications that require a direct correspondence between input sampling and digital output code are very limited. In most cases it is the signal waveform that is of interest and the digital filtering in the decimator only removes unwanted out-of-band noise.

Sigma-Delta ADCs and Other ADC Architectures

The most common ADC architectures are the pipeline, the parallel, successive-approximation (SAR) and sigma-delta. Each has its own performance space determined by effective resolution and signal bandwidth. Effective resolution is normally represented as effective number of bit (ENOB) and is obtained from SNR measurements using equation [1].

Figure 9 shows the performance space for these ADC converters. The highest bandwidth range is dominated by parallel ADCs, while SAR ADCs cover the low bandwidth and low-resolution spaces. The middle section is shared by sigma-delta and pipeline ADCs; sigma-delta reaches up to 10 MHz on the higher resolution range and pipelines extend beyond 100 MHz. The sigma-delta space overlaps the SAR and pipeline ADCs space at its lower resolution range.

The best architectural choice can be made by using system considerations covered in the previous section.

A sigma-delta ADC can be very easily reconfigured for different resolutions and bandwidths. Other ADC architectures only allow changing the clock frequency, adapting the current consumption to accommodate different signal bandwidths. For example, for half the signal bandwidth, the clock frequency is halved and the current consumption is halved. But sigma-delta ADCs have in addition, the option to work in a different oversampling rate by just programming the digital filter settings. That allows both signal bandwidth and SNR programming.
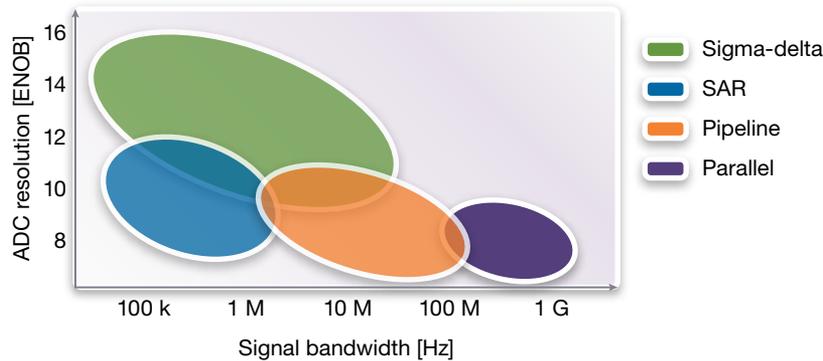

Figure 9: Performance space of most common ADC architectures

Application Examples Using Sigma-Delta ADCs

An example of a cellular radio receiver application using the classic direct conversion architecture is shown in Figure 10.

The signal received from the antenna is first amplified by a low noise amplifier (LNA) and then down converted to base-band I and Q paths. Other receiver architectures use additional stages of frequency conversion to bring the signal to the base-band, however the base-band chain itself is conceptually identical across the architectures. Each base-band path is composed of a filter to remove the out-of-band channels and interferers, an amplifier and an ADC. To accommodate the wide variation in received signal strength, the LNA and base-band amplifiers have programmable gain.

The signal can be very weak and surrounded by strong out-of-band interferers. These must be sufficiently attenuated in the filter so that the amplifier can be set to high gain to fit the signal to the full input range of the ADC. Since the resulting filter order would be too high, only partial attenuation is implemented in the analog filters. The ADCs must be sampled at 2X or 4X the Nyquist rate and have some additional dynamic range to accommodate the insufficiently attenuated interferers. A digital filter with sharp frequency response removes the residual interferers energy after the ADCs.

By allowing high oversampling ratios, the sigma-delta modulators allow further simplification of the analog filter requirements.

This option requires a higher dynamic range on the sigma-delta ADC to accommodate the lower signal amplitude.

Other application example is for audio and biomedical signals. These signals come from a microphone or other types of sensor and capture a very broad band that often extends well beyond the desired analysis band. For example, nature sounds may have significant energy above the audio range of 20 to 20000Hz, namely from some birds. And physiological measurements such as electroencephalography (EEG) and electrocardiography (ECG) are affected by body movement and broadband noise that can lead to important signal energy above the band of interest. These are, therefore, similar situations as those described above for wireless systems and benefit similarly of sigma-delta ADC solutions that allow moving the filters to the digital domain.
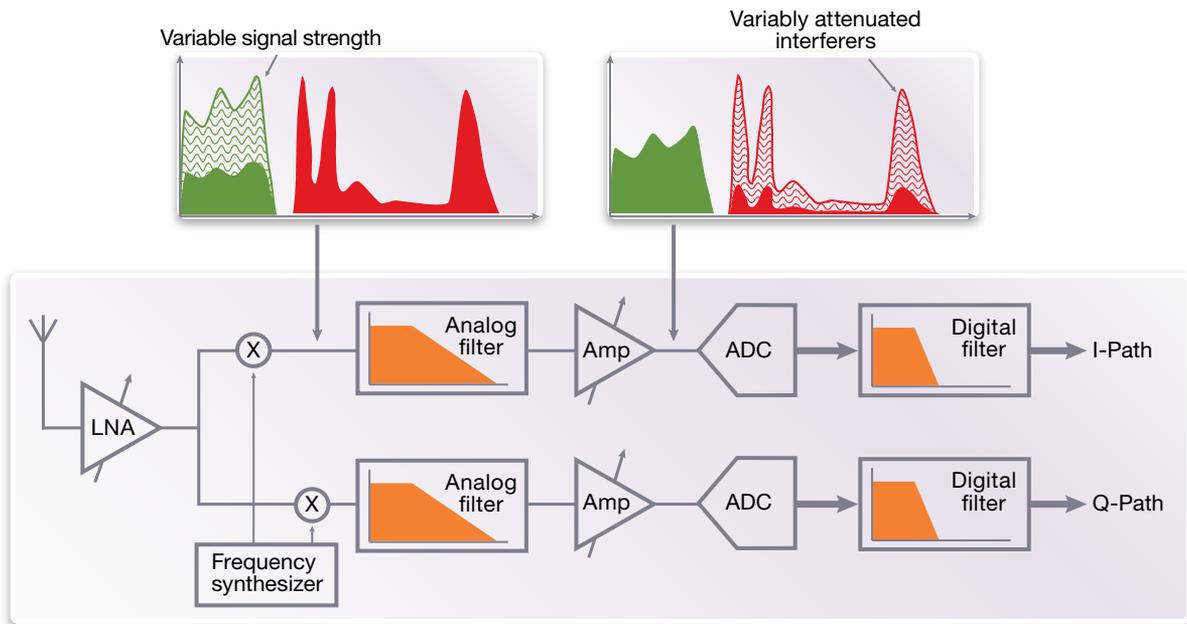
Figure 10: Radio receiver using Direct Conversion architecture

Sigma-delta ADCs are available from Synopsys featuring programmable bandwidth and dynamic range that can support a multitude of telecom standards, from 2G to LTE.

Due to its configurable nature, the sigma-delta ADC can be very efficient for systems where narrow band standards (e. g. 2G) and broadband standards (LTE) need to be processed in the same hardware.

GSM/EDGE is a 2G standard with a bandwidth of just 100 kHz in the base-band, but requires a relatively large dynamic range due to the high level of the adjacent channel interferers. The typically required dynamic range on the ADC is 70-80 dB. This can typically be achieved with an oversampling ratio of 130X, leading to an over-sampling frequency of 26 MHz (100 kHz * 2 * 130). Due to the relatively low clock frequency, the power consumption can be lowered.

LTE is an emerging 4G standard with multiple modes. The fastest mode requires a base-band bandwidth of 10 MHz. The dynamic range requirements are more relaxed at typically 60-65 dB, achievable with a modest oversampling ratio of about 10X and an oversampling frequency in the range of 200 MHz. At this high clock frequency the consumption must be increased.

A programmable ADC would normally support different combinations of intermediate dynamic range and signal bandwidths between the above given examples. Such ADCs are typically available as a matched pair, one for the I path and the other for the Q path. Synopsys provides this intellectual property (IP) core complete with clock generation, biasing, references and CIC digital decimating filter.

Synopsys IP offering includes also lower frequency sigma-delta ADCs that are specifically designed for sensor signals, such as EEG, ECG, accelerometers, pressure and stress. These typically require very flexible front-ends supporting both single-ended and differential formats, with programmable gain and capability for multiplexing multiple inputs. Also popular in microcontroller applications that include programmable interfaces to external sensors and transducers requiring more resolution than the typical 12-bit that SAR ADC can provide.

Conclusion

Sigma-delta ADCs are ideally suited for integration into advanced SoCs because they take full advantage of the high-speed and high density of devices available in deep submicron process technologies. Sigma-delta ADCs offer flexibility for programming the bandwidth and resolution, thanks in part to a wide range possible in both the sampling frequency and oversampling ratio. Their high oversampling rate also enables the simplification of analog filtering in front of the ADC. By being aware of the benefits of the sigma-delta ADCs, SoC designers are better equipped to address the demanding requirements of their applications and produce superior solutions that deliver the performance, area and power gains.

With more than 15 years of experience in developing data conversion IP solutions, Synopsys offers a comprehensive portfolio of over 200 silicon-proven DesignWare® Data Conversion IP solutions consisting of oversampling sigma-delta ADCs, pipeline ADCs, SAR ADCs, current-steering DACs and much more. The DesignWare Data Converter IP products offer very low power dissipation, small area and support a wide range of process technologies and foundries ranging from 180-nm down to 28-nm.

For more information on the Synopsys DesignWare Analog IP Data Converters portfolio, visit: http://www.synopsys.com/dataconverters