

Putting Continuous Time Delta-Sigma Modulators to Work

Providing a route to the lowest power, high dynamic range ADCs and offering new levels of integration.

Contents

1	Data conversion fundamentals	2
1.1	Contemporary pipeline converter overview	2
1.2	Some downsides of the pipeline ADC	3
1.3	Historical delta-sigma conversion	4
2	Introduction to delta-sigma conversion	4
2.1	Over-sampling allows sample frequency/SNR trade-off	4
2.2	Noise shaping	5
2.3	Digital filtering and decimation	6
3	Continuous time versus discrete time systems	6
3.1	A universal problem with discrete time sampling	7
3.2	The continuous time implementation	8
4	Advantages of Xignal's CT$\Delta\Sigma$ implementation	8
4.1	Ultra-low power consumption	8
4.2	Inherent anti-alias filter simplifies ADC input circuit	9
4.3	Easy to drive input stage	10
4.4	CT $\Delta\Sigma$ modulator and digital filtering network	10
4.5	Low jitter PLL provides an accurate sample clock	10
5	Xignal's enabling technologies	11
6	Conclusions	11
	Acknowledgement	12

Continuous time delta-sigma ($\text{CT}\Delta\Sigma$) analog-to-digital conversion technology shatters the conventional wisdom that pipeline ADCs are the only conversion technique available for high-speed, high dynamic range applications. $\text{CT}\Delta\Sigma$ technology offers lower power operation, better dynamic performance, economy of design and specific functional benefits.

In short, $\text{CT}\Delta\Sigma$ technology means:

1. An inherently power-efficient architecture that eliminates power-hungry sample-and-hold amplifiers (SHA) and the wide-bandwidth gain stages essential to the pipeline ADC concept.
2. An alias-free Nyquist sample range that is made available by exploiting inherent over-sampling and on-chip digital filtering. Digital filtering permits tailoring of group delay performance and the signal transfer function to specific applications.
3. Simplicity of application through the integration of a clock and low jitter PLL, the elimination of anti-aliasing filters and the integration of input gain stages to simplify input signal path design of high-resolution data conversion systems.
4. Switchless design future-proofs $\text{CT}\Delta\Sigma$ technology, through its easy migration to the next generation of CMOS processes, enabling the speed and power benefits of the process roadmap to be attained.

Signal Technologies' $\text{CT}\Delta\Sigma$ architecture supports **high-resolution analog-to-digital conversion systems from 10 to 16-bit and beyond with sampling rates up to 100 MHz.**

1 Data conversion fundamentals

In applying a sampled data system, no treatment is complete without an understanding of Nyquist sampling theory. This has significant implications for system design. It is also a useful introduction to the concept of over-sampling at the heart of delta-sigma conversion. By general implication, if Nyquist's theorem is applied, then a representation of an input signal can always be reconstructed from the time invariant samples produced by an analog-to-digital converter.

Over the years, to perform such signal acquisition, a variety of techniques have been proposed. An early analog-to-digital conversion technique was the direct conversion or flash technique. The flash converter features a cascade of parallel comparators connected to a resistive, divider-derived voltage reference. The reference voltage for a given comparator is designed to be 1 Least Significant Bit (LSB) greater than the neighboring comparator one step below it. The resulting output from each comparator is the so-called thermometer code, which can be converted into a binary-weighted digital output. For an n-bit resolution flash ADC, $2^n - 1$ comparators are required, consuming a considerable power budget. Also the challenge of providing a precise reference supply to each comparator, and ensuring that each comparator closely matches its neighbors, further complicates the design. From a chip design perspective, a large investment in die real estate is required. For all practical purposes, flash converters are normally limited to 8-bit resolution (255 comparators). Most subsequent design effort to reduce power in high-speed conversion has worked to minimize the number of comparators applied, for example by interpolating and folding designs.

1.1 Contemporary pipeline converter overview

Today, it is broadly assumed that the pipeline converter provides the highest sample rates whilst yielding high dynamic range. It is used as a standard in data conversion applications at 10-bit and higher resolutions and for sample rates from 5 MHz to 100 MHz or more. The architecture reduces the number of comparators needed by deploying multiple low-resolution flash conversion stages cascaded together to form the pipe. Although the resolution of each conversion stage is reduced, saving on comparators, the first stage must be designed with linearity at least as good as the maximum resolution of the ADC (12-bit linearity for a 12-bit ADC). Different pipeline implementations exist, but all work by reducing a multi-bit resolution conversion into several lower resolution "flashes" that are processed synchronously. At each stage in the pipe, a reconstruction of the previous stage's quantized output, generated within a DAC, is subtracted from the original input signal. The residual signal is then amplified prior to moving onto the following stage for ever finer quantization.

Basic Nyquist Sample Theory

The theory states that for an error free reconstruction of a signal, given time invariant sampling, the system must sample the input signal at a minimum rate of twice the input signal's bandwidth.

This constraint is set because the sampling process mixes the input signal with the sampling signal. It generates an infinite spectral series (harmonics) comprising paired multiples of components in the frequency domain. Sampling a signal f_{in} at a sample rate of f_s results in indistinguishable alias signals found as spectral spurs centered at $x(f_s \pm f_{in})$ – where x is an integer multiple.

The lowest input frequency spur ($f_s - f_{in}$) tends to be most important as it falls within the Nyquist bandwidth – the range of signal frequencies of interest to the sampling process. For this reason anti-alias filters are used to eliminate out of band signals.

For over-sampling converters, the sample clock is set considerably higher than Nyquist criteria demands, providing dynamic range improvements.

The pipeline conversion process requires several sample clock cycles before the final conversion result is available (this is the latency of the pipeline). During this time the sampled signal may well change and for this reason, the input must be sampled and held within the ADC to allow each stage to quantize the signal. This is the role of the sample-and-hold amplifier (SHA). This acquires the input signal and holds it to better than 0.5 LSB for the duration of the conversion. Once all sub-stages have a valid conversion result, a digital correction block constructs the final multi-bit result.

1.2 Some downsides of the pipeline ADC

The pipeline ADC is clearly capable of high dynamic performance levels. However, beyond 12-bits resolution there are major challenges in design including the sizing of capacitors and their ratios for charge storage during the conversion process. The maximum dynamic range of the ADC is dictated by the thermal noise occurring in the sample and charge transfer capacitors. Using larger capacitor values minimizes this. However as the sampled signal moves through the pipeline, transferring the charge associated with a given signal demands very high gain bandwidth circuits to ensure stage settling times fall within the limits set by the high frequency signals of interest that are being sampled. Doing this at high-speeds and resolution is no mean feat, especially with realizable capacitors.

To maintain linearity, schemes are needed to calibrate and correct for the limits in component matching achievable with

current process technologies. Because of the sensitivity of high performance analog circuits found in the pipeline design, it is challenging to migrate designs from one process to another. As operating voltages necessarily reduce from one process generation to the next, the input signal headroom is compressed. Furthermore, designing switches with greatly reduced threshold voltages that work well in deep sub-micron processes becomes increasingly challenging. Continuous time delta-sigma provides a chance for engineers to uncouple high dynamic range designs from these process limitations.

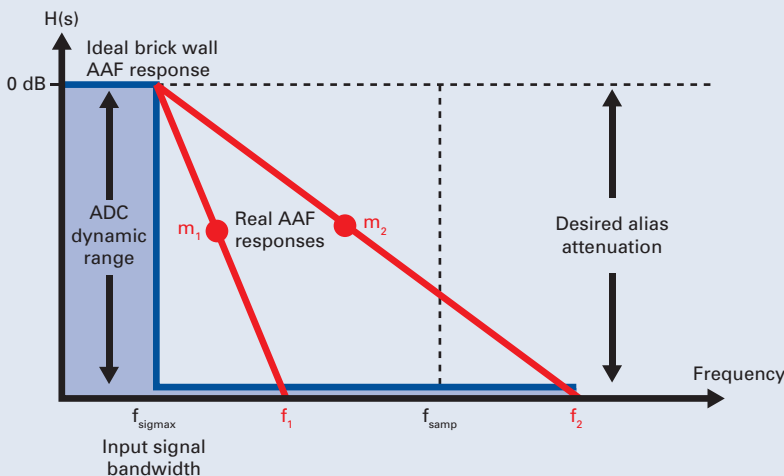
A final, often-underplayed issue with pipeline ADCs is that they form only part of a data conversion system, the complete design of which is left to the system designer. Challenges arising include sourcing a low jitter clock source and the careful design and selection of input stages including anti-alias filters.

ADC application notes implore designers to take care of the external environment in which the ADC functions. Consider the task of anti-alias filter (AAF) design. Steep filter attenuation characteristics are hard to achieve, tempting a designer to consider over-sampling the signal of interest, thus lowering the demands on the AAF.

This idea is expanded on in the sidebar 'Wasted Bandwidth'. This shows that over-sampling stretches the Nyquist zone, lowering demands on filter roll-off. However, the system implications of this method are increased system power and higher processing speeds demanded of the back-end DSP system. That's fine if

Wasted Bandwidth – Nyquist’s criteria and its real world implication

Practical implementations of anti-alias filters (AAFs) have a direct impact on the design of sampling systems. To eliminate any alias signals falling into the sampling bandwidth, the AAF must achieve an attenuation of out of band signals so that they fall below the quantization level of the system. Nyquist’s theorem is re-stated as equation 1 below. A brick wall low pass filter is ideal for the AAF but cannot be realistically achieved. Thus, for a standard Nyquist sampling system, the AAF stretches the signal bandwidth that must be digitized as shown by frequencies f_1 or f_2 in the diagram below (representing filters of order m_1 & m_2). Applying Nyquist’s theorem in the real world means that the sample rate is set based on the intercept of the AAF roll-off with the system noise floor. The intercept is a function of desired attenuation (which is related to the SNR or the system) and the order of the AAF (or roll-off rate). Thus non-ideal roll-off leads to wasted system bandwidth. A modified version of Nyquist’s theorem is thus proposed as equation 2, which captures this effect. A key advantage of the CT $\Delta\Sigma$ architecture is that with over-sampling the CT $\Delta\Sigma$ removes the external AAF filter requirement providing an AAF-free sample range.



Equation 1
 $f_{\text{samp}} = 2 \times f_{\text{sigmax}}$ – per Nyquist

Equation 2
 $f_{\text{samp}} = 2 \times f_{\text{sigmax}} \times \frac{A(f)}{m}$

Where A = attenuation at a given frequency and m is the order of the AAF.

— Wasted system bandwidth.

your business is selling DSP horsepower. In the real world, applying Nyquist's criteria to system design is better expressed by a revised relationship that captures the bandwidth wastage of the anti-alias filter. As will be shown, the continuous time delta-sigma converter provides a true anti-alias filter-free Nyquist sample range.

1.3 Historical delta-sigma conversion

Delta-sigma ($\Delta\Sigma$) conversion has been applied for many years. The $\Delta\Sigma$ converter uses a low-resolution quantizer - often only 1-bit - clocked at rates considerably greater than Nyquist demands. The quantizer creates a large number of low-resolution samples that, averaged over time, yield an increased dynamic range.

The traditional $\Delta\Sigma$ design is a radical departure from analog topologies like the pipeline ADC. For the $\Delta\Sigma$ ADC, the analog design is potentially straightforward given the exceptional linearity of a 1-bit (2 level) quantizer. Its complexity moves more into the digital domain where filtering and decimation is needed to reconstruct output data and remove out-of-band noise.

The $\Delta\Sigma$ technique has traditionally been deployed for relatively low frequency tasks including audio and precision measurement systems. High speeds have previously been hard to realize for several reasons that will be discussed further. Before considering the specific design of CT $\Delta\Sigma$ architecture, it's worth building an understanding of the basic operating principles behind the delta-sigma technique. To understand this, one must build an appreciation of over-sampling, noise shaping, digital filtering and decimation.

2 Introduction to delta-sigma conversion

The $\Delta\Sigma$ converter operates at sample rates beyond that dictated by Nyquist's criteria. To ease understanding of the delta-sigma modulator, Figure 1 provides a high level view of the simplest single order $\Delta\Sigma$ modulator block at the heart of the concept. Its design simplicity is clear on first glance. It comprises a

summing node and an integrator connected to a comparator. The comparator's output provides the input to a 1-bit feedback DAC that then closes the modulator's feedback loop.

The modulator compares the input signal against a voltage reference level fed back from the DAC. The comparator is clocked at the over-sampling frequency. As with all feedback systems, assuming enough loop gain is applied, then the modulator output must be equivalent to the input signal at equilibrium. The resultant two level output will be a pulse stream. The density of '1s' or '0s' forming this pulse stream is a direct digital representation of the input signal, the DAC switches between $\pm V_{ref}$ to close the control loop.

In common with all other conversion techniques, quantization errors arise within the modulator ultimately limiting dynamic range. Quantization error is a measure of an n-bit converter's failure to represent precisely the analog signal in the digital domain.

For the $\Delta\Sigma$ ADC, resolution increases are gained by balancing three design aspects: the over-sampling ratio, $\Delta\Sigma$ modulator order, and quantizer resolution. By careful consideration of these variables a number of important performance trade-offs are possible.

2.1 Over-sampling allows sample frequency/SNR trade-off

Over-sampling improves dynamic range (DR). This effect is best considered by exploring quantization noise using a simplistic graphical treatment. Figure 2, shows how the quantization noise power in the Nyquist bandwidth ($f_s/2$), when over-sampled, spreads over the wider $Mf_s/2$ bandwidth, effectively lowering the noise floor. It can be shown that this improvement is directly related to the over-sampling ratio (M) by the following relationship - true for a single order modulator:

$SNR = 10 \log_{10} M$, where $M = f_{os}/f_s$ and f_{os} is the over-sampling frequency.

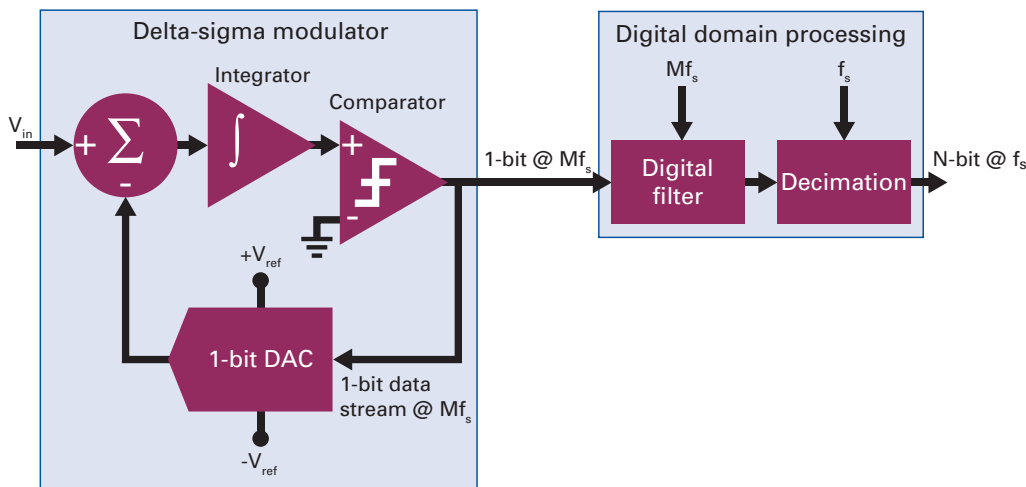


Figure 1: Simple 1-bit $\Delta\Sigma$ modulator block diagram.

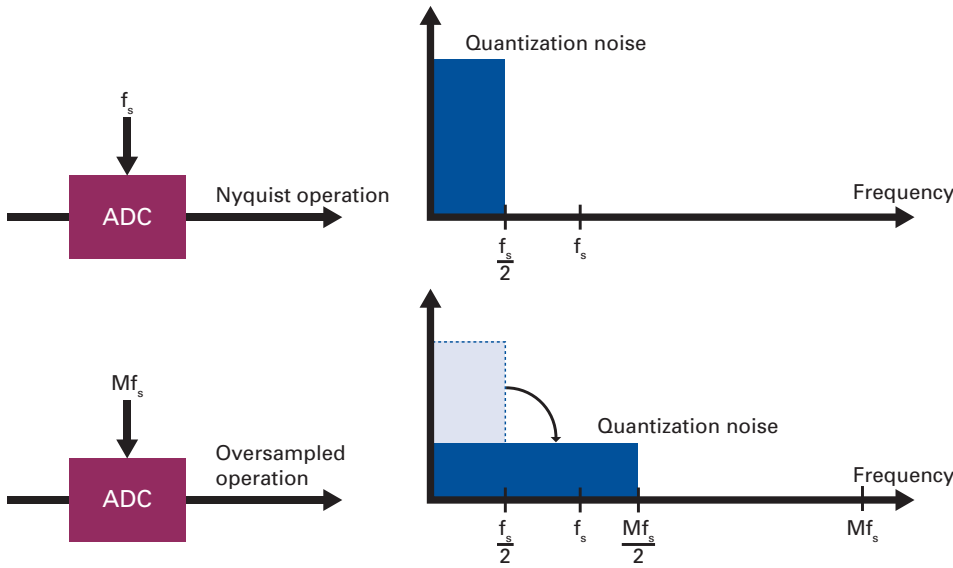


Figure 2: Over-sampling spreads quantization noise energy lowering the effective noise floor.

This gives an SNR improvement of 6 dB (or 1-bit's worth) for every quadrupling of the sample rate (SNR improves by 4^{n-1} , where n is the resolution). This is a fairly costly and impractical way of improving resolution. For example, consider sampling a 5 MHz bandwidth signal and aiming to achieve a 12-bit dynamic range. To achieve the 74 dB ($6 \times 12 + 1.76$) the SNR requires an over-sampling frequency of 21 THz ($4^{11} \times 5 \times 10^6$). DR is therefore more easily increased by two alternative methods. Firstly, increasing the resolution of the quantizer offers an immediate enlarged DR. Secondly, DR gains can be achieved by adding additional integration stages to the modulator.

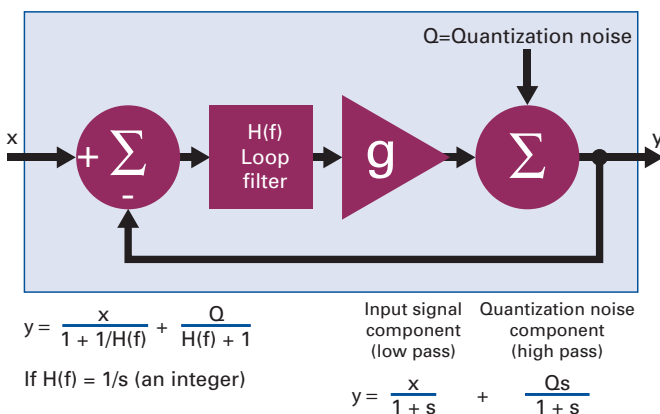
2.2 Noise shaping

Noise shaping is a property of $\Delta\Sigma$ ADCs resulting from the application of feedback that extends dynamic range. This feature is best illustrated by the mathematical analysis of the feedback control loop of the $\Delta\Sigma$ modulator as modeled in the frequency domain, illustrated in Figures 3a.

This model reveals the key value of the $\Delta\Sigma$ modulator. A closed loop modulator works as a high-pass filter to quantization noise and as a low-pass filter to the input signal. A first order

modulator produces a characteristic shown in Figure 3b. The net effect of this is a further increase in dynamic range of 9 dB for each doubling of the sample rate. Additional integrator stages may be added to the modulator loop and can increase the dynamic range still further as shown in Figure 3c. A detailed treatment of multiple order feedback in the modulator is a complex subject as it can have significant implications with regard to stability analysis and system performance. However, it can be shown that additional integrators within the loop can increase the steepness of the noise characteristic and thus gain further dynamic range increases.

Figure 4 shows simulation results for the noise power density for the $\Delta\Sigma$ modulator used in Xignal's ADC. This FFT plot (with 65k points) of the modulator illustrates the noise power density (per FFT bin) relative to the input signal frequency. The simulation was driven with an input signal frequency of approximately 4.8 MHz. The minimum noise power density achieved by this modulator is 166 dBc/Hz (in the pass band). Note the characteristic of the out of band noise i.e. those frequencies above 20 MHz. Here noise power levels rise at the rate of 21 dB/octave, a tell-tale sign of a third order modulator. Actually deployed within this design is a



x is the input signal & y the $\Delta\Sigma$ modulator output

The quantization error is modelled as Q added to the modulator output.

Assuming $g = 1$, and the transfer function is $1/f$ then the resultant output can be seen in terms of an input signal component and a noise component.

Figure 3a: Frequency domain model of $\Delta\Sigma$ modulator.

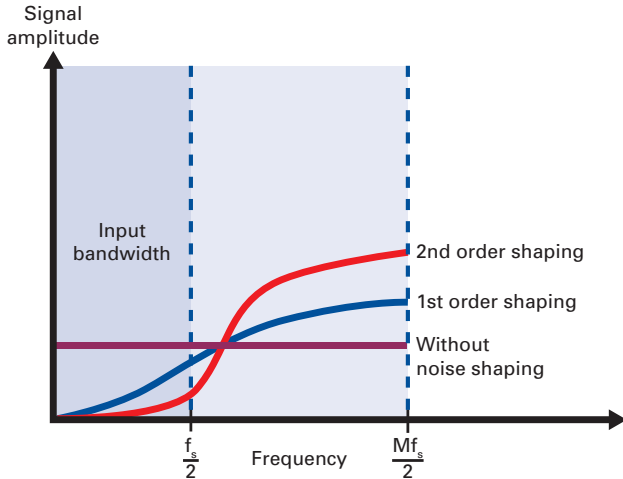


Figure 3b: 1st & 2nd order noise shaping characteristic.

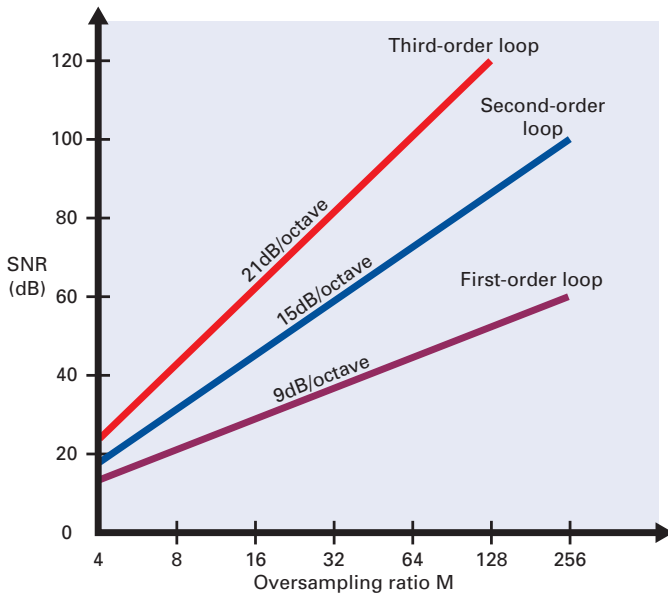


Figure 3c: SNR versus oversampling ratio for single & multi-order loops.

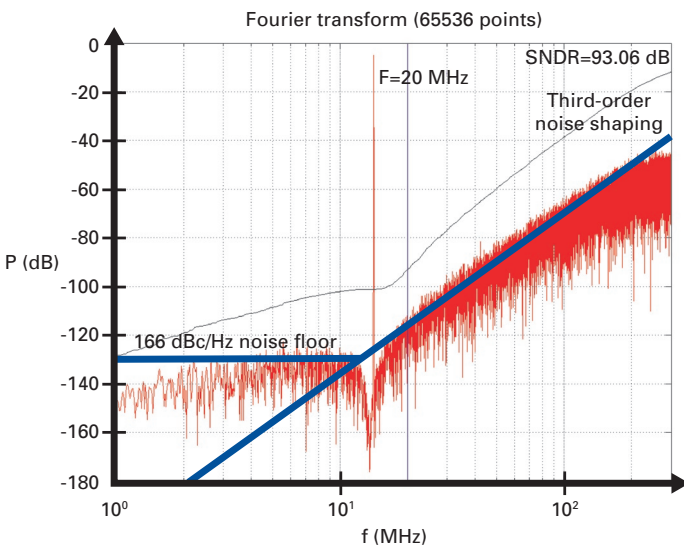


Figure 4: Simulation of noise transfer function in CTΔΣ ADC.

16 levels, or 4-bit, quantizer that delivers 14-bit dynamic range at modest over-sampling rates. Having established a modulator system capable of achieving these low noise levels, the next stage is to apply filtering to eliminate out-of-band noise, and decimation to re-sample the data.

2.3 Digital filtering and decimation

Digital filtering is a complex subject. The digital filter must reject all signal components within the serial data stream that occur beyond the Nyquist bandwidth. Simplistically, two frequency selective filter structures can be implemented in the digital domain. They are the finite and infinite impulse response filter systems (FIR and IIR). These are widely described in many public domain articles. Because of their design ease, FIRs are more widely used and they offer the benefit of linear phase response.

IIR filter design is more complicated by virtue of the feedback included. Deriving filter coefficients from a target transfer function can be quite complex. The potentially infinite response of the IIR filter means there is always a possibility for the filter to become unstable. In addition, group delay, the rate of change in phase with angular frequency, can become quite significant. In many systems group delay can have direct adverse effects on performance.

Many degrees of freedom exist for optimization of the signal transfer function in a CTΔΣ ADC through combining different filter algorithms, however optimal solutions may require many cascaded stages of FIR and IIR sections. Digital filtering allows for the data reduction or down-sampling necessary to provide output data at the originally intended sample rate. Note that over-sampling provides large amounts of redundant data. The process of sample rate reduction is called decimation.

All the basic elements of a delta-sigma ADC have now been presented. In Figure 5, the signal flow within the blocks of a ΔΣ ADC is illustrated to put the three major operating principles in context. Firstly, over-sampling spreads quantization noise. Secondly, noise shaping reduces the in-band noise at the expense of higher out-of-band noise. Thirdly, digital filtering attenuates out-of-band noise and signal components.

3 Continuous time versus discrete time systems

The majority of pre-existing ΔΣ converters found in audio and precision applications exploit switched capacitor, discrete time (DT), loop filters within the modulator for noise shaping. Understanding the behavior of discrete time switched capacitor filters (SCFs) relies on the application of sample data theory. SCFs by their sampling nature create mixing products of their own, so care must be taken to control unwanted out-of-band noise folding back into base-band. The DT ΔΣ design is inherently susceptible to this type of noise aliasing. Even so, the advantage of DT schemes is their relatively simple architecture, the way that increased sample rates produce dynamic range improvements, and their compatibility with VLSI CMOS processes.




Figure 5: Putting noise shaping & digital filtering together.

With switched capacitor filters, pole and zero settings are achieved by selection of capacitor ratios and sampling clock frequency. The capacitors must be relatively large for maximum SNR. This is because capacitor thermal noise is inversely proportional to capacitance. Fortunately, the DT filter characteristic scales with the clock frequency, ensuring that filter performance matches the sample clock rate at any given time. However, as the clock frequency increases, the dynamic power consumed increases.

In the last few years, with the ease of embedding them into System-on-Chip (SoC) designs and the inherent low power, attention has turned to the broader application of $\Delta\Sigma$ modulators. Switched capacitor stages act as a limit on the maximum signal bandwidth in the DT $\Delta\Sigma$ ADC. Moving to continuous time (CT) loop filters raises the bar and opens up a raft of new application possibilities. These include wide base-band sampling out to several tens of megahertz to under-sampling RF signals in band-pass designs.

3.1 A universal problem with discrete time sampling

So far, the operation of both the pipeline and DT $\Delta\Sigma$ ADCs has been considered. It may not be so obvious that a common

design thread ties them together. In discrete time, sampling an input signal requires that the signal be acquired at a precise moment in time. For an accurate representation of the input signal to be acquired on a hold capacitor, it is necessary that the input stages settle to a finite level in a time period driven by the system sample rate needs. This level is dictated by the accuracy limits of the system (e.g. for a 12-bit system settling to 244 ppn). Settling relies on the exponential characteristic of charge transfer within a simplified resistor & capacitor (RC) network. During the acquisition phase of this circuit its performance is dominated by the exponential time constant τ (tau) of the system. The number of time constants for an RC system to settle can be derived mathematically and shown to be $\tau = 8.3 \tau_c$ to achieve 12-bit and $\tau = 9.7 \tau_c$ for 14-bit settling. In practice the total time is more than this as it is also dependent on the slew rate time t_{sr} as shown in figure 6.

This settling time eats into the sample time period of the system. At 40 MSPS a conversion system will have a sample period of just 25 ns, which sets the maximum time limits for circuit settling. At 14-bits resolution, τ_c is a maximum of only 2.5 ns (τ approx 10). This drives a need for very high gain

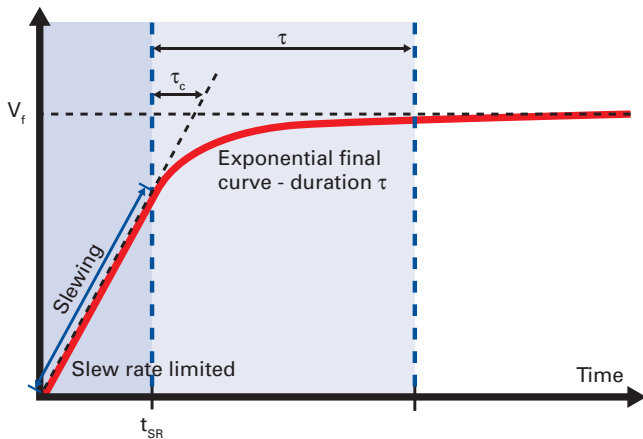


Figure 6: Settling time limitation of discrete time sampled systems.

bandwidth circuits within the acquisition signal path. In fact, this consideration shows that a converter system, be it a pipeline or a $DT\Delta\Sigma$ architecture, must be designed with circuits that work with bandwidths many times that of the input signal. Discrete time circuits therefore have to burn excess power to process a given bandwidth.

The move to a continuous time strategy eliminates the settling time issue altogether. Thus, allowing either a lower power $CT\Delta\Sigma$ ADC implementation, versus discrete time, at a given sample rate or a higher sample rate for a given power budget.

3.2 The continuous time implementation

In continuous time there is no acquisition phase, meaning a high performance sample-and-hold stage can be eliminated. Continuous time does not require the high gain bandwidth stages necessary to force rapid settling, meaning power in these stages is reduced. Albeit, a little more power will be demanded from the post modulator digital filter stages. The CT loop filter reduces the power demands of the modulator and eliminates the additional discrete time sampling effects seen in the $DT\Delta\Sigma$ ADC. There is no down-mixing of noise, eliminating any additional spectral spurs in the base-band. Continuous time filters, though more complex to implement, are far easier to drive than DT SCFs, as the capacitive load creates current pulses, demanding some input buffering yet another power consuming element of discrete time systems.

The downside of $CT\Delta\Sigma$ is that the ADC designer’s task is more complex. For example, consider that the filter characteristic is no longer tied to a filter sample clock, as with SCFs. Now, the filter performance is dependant on conventional active filter design rules. If sample rate is changed to match input signal bandwidth the CT filter must be tuned. Therefore, a potential limit on the $CT\Delta\Sigma$ implementation is how to ensure a wide range of sample rates can be supported from a single product platform. At Xignal this problem is solved using adaptive filter networks in combination with calibration techniques.

The circuit design needs in $CT\Delta\Sigma$ are tough. For high-resolution implementations the loop filter must have significant gain if high

An ADC Figure of Merit (FOM)

The analog program committee of the IEEE International Solid-State Circuits Conference suggested an objective figure of merit (FOM) for ADC comparisons that takes into account power consumption, effective resolution and sampling rate. With units of energy it represents the energy used per conversion, and is calculated as follows:

$$FOM = \frac{P}{(2^{ENOB} \times f_s)}$$

P is the power dissipation in Watts, ENOB is the effective number of bits at a sample rate f_s in Hertz. FOM yields a result in Joules/conversion.

Rarely found in manufacturer’s data, the FOM has merit in comparing the performance of a particular ADC product or specific architecture. Its value lies in identifying best energy efficiency compared to raw dynamic performance of the converter. In using the ENOB specification it quantifies all noise and distortion sources introduced by the converter.

linearity is to be achieved. This places very tight demands on building compound gain stages, especially within a mainstream sub-micron CMOS process. Xignal has developed sophisticated multi-path and cascaded gain stages that, whilst operating at a mere 1.2 V, achieve 80 dB of gain for a 30 MHz bandwidth. This is a significant achievement in 0.13µm CMOS.

Cracking these design challenges opens up the possibility of a low power alternative to the pipeline converter. In fact the $CT\Delta\Sigma$ implemented by Xignal is a lot more than an integrated ADC core. This offers ease-of-use beyond the conventional pipeline ADC. Xignal’s advanced technology provides a complete data conversion system, designed to operate seamlessly over a wide range of sample rates. **This is arguably the first ADC that performs close to the ideal of a true Nyquist converter with no high performance, expensive, external components.**

4 Advantages of Xignal’s $CT\Delta\Sigma$ implementation

Xignal Technologies’ implementation of a continuous time $\Delta\Sigma$ quantizer and digital filtering has lead to an innovative ADC core supporting resolutions up to and beyond 14-bits. The architecture offers considerable flexibility and enhanced integration options. In its first development, Xignal’s general purpose ADC core is combined with several features designed to simplify their application in a number of data conversion roles. The benefits of this implementation are now considered further.

4.1 Ultra-low power consumption

The most significant benefit for the $CT\Delta\Sigma$ architecture is its power consumption. **A common way to compare the performance of ADCs in absolute terms is the energy figure of merit (FOM).** This measure is explained in the sidebar and Table 1 shows typical FOMs of a number of products.

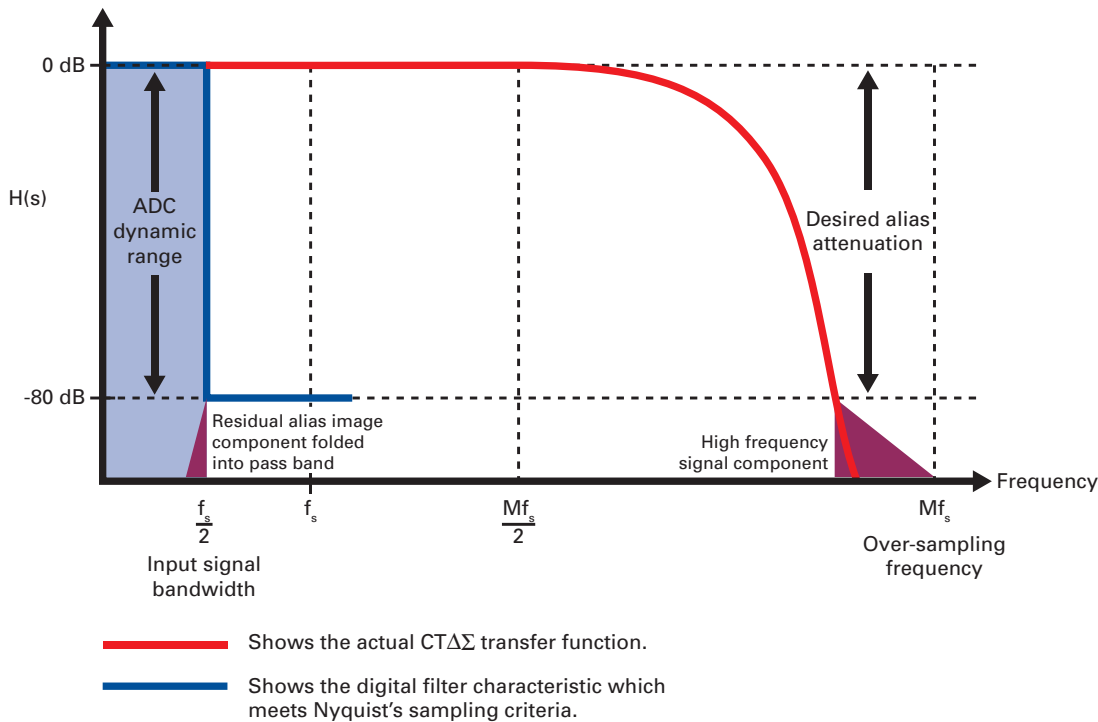
Converter Topology	Resolution	Sample rate	Typical FOM
1. Flash	Typically up to 8-bits	Multi-GHz	VP1058 – Zarlink, 8b, 25 MSPS ADC Status: Obsolete (included for comparison only) FOM = 259 pJ/conv
2. Folding/Interpolating	6 to 10-bits	Up to 1 GHz	ADC081000 – National Semiconductor, 8b, 1000 MSPS Status: Production, FOM = 13.1 pJ/conv
3. Pipeline	10 to 16-bits	5 to 150 MHz	ADS5421 – TI, 14b, 40 MSPS Status: Production, FOM = 5.5 pJ/conv AD9244-40 – Analog Devices, 14b, 40 MSPS Status: Production, FOM = 2.78 pJ/conv MAX1260 – Maxim, 12b, 40 MSPS Status: Production, FOM = 2.54 pJ/conv LT2247 – LTC, 14b, 40 MSPS Status: Production, FOM = 1.05 pJ/conv
4. Conventional $\Delta\Sigma$ (discrete time)	>16-bits	200kHz to 1MHz	ADS1605 – Texas Instruments, 16b, 5 MSPS ADC Status: Production, FOM = 10.7 pJ/conv
5. Advanced CT $\Delta\Sigma$	10 to 16-bits	1 to 100 MHz	XT11 development – Signal Technologies, 14b, 40 MSPS Status: In development , FOM = 0.4 pJ/conv with PLL & CLK,

Table 1.

Simulation of Xignal's latest power optimized 14-bit, 40 MSPS ADC shows it will provide a 76 dB SNR (at $f_{in} = 6$ MHz) and a power consumption of only 70 mW. This gives a calculated FOM of 0.4 pJ/conv, some 50% lower than the best existing pipeline converters. Note that this includes the power over for the integrated clock and PLL, underscoring the true potential of this architecture.

4.2 Inherent anti-alias filter simplifies ADC input circuit

The need for anti-aliasing filters demanded by Nyquist's theorem has been previously discussed. It was shown that the $\Delta\Sigma$ ADC architecture eliminates complex input filtering due to the effect of over-sampling in combination with digital filtering, as shown in Figure 7. This feature cannot be over stated. In a 14-bit system


Figure 7: Oversampling and its effect on anti-alias filtering.

offering a theoretical SNR of 86 dB ($SNR = 6.02 \times n + 1.76$), the filter needed is dependent upon the application and can be a significant design problem. The design of an analog anti-alias filter capable of steep cut-off characteristic and gain flatness in the pass band is a challenging task demanding high order filter networks. Furthermore, at high frequencies the only option is to use passive circuit elements. Alternatively, the sample rate may be increased, wasting sampling bandwidth to ease the AAF design. In Signal's CT $\Delta\Sigma$ ADC the anti-alias filter is inherent to the design. With filter performance characteristics set in the digital domain, a very high level of pass-band flatness and steep roll-off is possible. The current digital filter allows 90% of the available Nyquist bandwidth to be exploited whilst offering a pass band ripple of only ± 0.0002 dB, and an 80 dB stop-band attenuation. Group delay for this filter is only 0.33 samples.

4.3 Easy to drive input stage

It's not only AAFs that CT $\Delta\Sigma$ technology eliminates. Removing the sample capacitor found in discrete time designs has an important consequence on the input stage of the ADC. As shown previously, slew and settling constraints force high bandwidth components to be used in discrete time implementations. Adopting a continuous time system however creates a considerably simplified input structure for the ADC. In fact the equivalent input circuit is modelled by a resistor driving into a virtual earth. Using this current driven input stage has the effect that the input voltage range is no longer limited to the extent of the supply voltage range of the convertor.

Xignal's first product achieves a $4 V_{p-p}$ input range despite operating from a single 1.2 V supply. This feature not only simplifies connecting real-world signals to the CT $\Delta\Sigma$ ADC, it also means that no extra discrete high bandwidth differential input driver IC is needed. Such a circuit is mandated for pipeline ADCs to handle the complex capacitive input impedance and

the resultant sample glitches. This further emphasises CT $\Delta\Sigma$'s system level advantages of power consumption, component reduction and performance.

4.4 CT $\Delta\Sigma$ modulator and digital filtering network

Xignal's new product uses a third-order CT $\Delta\Sigma$ modulator, designed around a 4-bits quantizer stage helping it to achieve considerable dynamic range at a relatively low over-sampling rate of 16. The fully differential input signal path has a bandwidth of 30 MHz. The internal sample clock operates at 640 MHz. The base technology today allows for increased sample rates to 80 MSPS (at 14-bits) with an over-sampling clock rate of approximately 1.3 GHz.

The resulting data stream is then fed to a cascaded multi-stage digital filtering system. Its design aims to optimize phase and amplitude performance whilst offering a compact low power implementation. Included within this filter is a user selectable equalization stage that reduces total group ripple variation to 0.33 samples. One of the design achievements for this product was creating a part able to operate over a broad range of sample rates. Through self-adaptive tunable loop filter components the ADC is optimized for sample rates from 20 to 40 MSPS.

4.5 Low jitter PLL provides an accurate sample clock

Because of the over-sample rates of $\Delta\Sigma$ ADCs, it can be imagined that the sample clock, and in particular its jitter, plays an important role in establishing the dynamic performance of the system. In almost all past $\Delta\Sigma$ implementations, selection of a clock source is left to the designer's discretion. Xignal integrates a high performance clock source on-chip with the ADC core. All that this ADC requires is a low cost crystal, parallel connected to its clock input. The clock is connected to a high performance PLL block that uses an on-chip LC tuned circuit to create a high-Q resonator, thereby creating a very precise clock source.

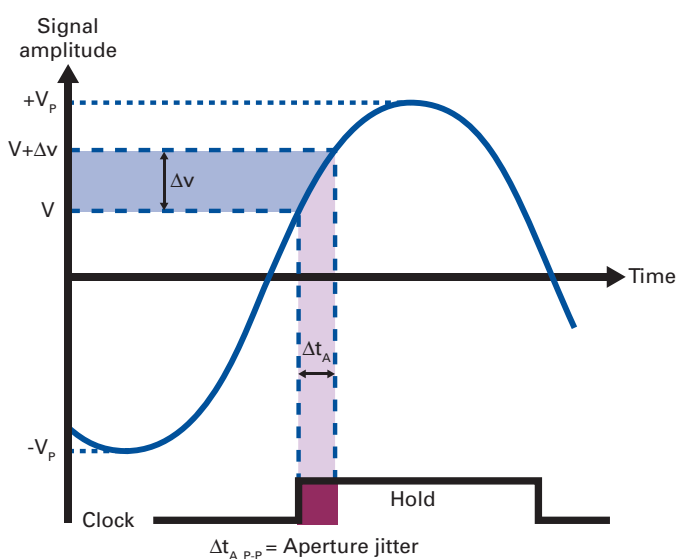


Figure 8: Clock jitter impact on SNR.

The aperture error is less than 1 LSB, if:

$$\Delta t_A < \frac{1}{2^n \times \pi \times f_{\max}}$$

where n is the resolution of the ADC and f_{\max} is the maximum signal frequency

Alternatively, an external clock can drive the ADC. High frequency jitter from an external distributed clock tree will be removed provided its jitter falls outside the 350 kHz PLL bandwidth of the jitter cleaner circuit. However, a further advantage of the on-chip precision clock is that it can be routed to external circuits and used as a system reference clock for other time-critical parts of the system, potentially eliminating the extra cost of a low jitter source, saving both design effort and board area.

A low jitter clock is a crucial function in all high-speed, high-resolution data conversion systems. Phase accuracy of the sample clock has a major impact on measured performance. In fact, decibels of dynamic range are easily sacrificed by picoseconds of phase jitter. Figure 8 shows the mathematical derivation of maximal clock jitter for a given resolution and input signal frequency. For a 10 MHz bandwidth signal, at 12-bit resolution, clock jitter must be less than 3 ps RMS. For 14-bits, this demand drops to 1 ps RMS. Further increasing input frequency or resolution clearly adds to the jitter problem. Xignal's on-board PLL generates a sample clock with just 300 fs ($f = \text{femto} = 10^{-15}$) RMS jitter. This helps define the sample point very precisely and minimizes any resultant spectral spreading.

5 Xignal's enabling technologies

Why, given the extensive research into $\Delta\Sigma$ converters over the last 50 years, has it taken so long to get to high-speed wide bandwidth implementations, especially given the relative advantages of the technology presented? The answer is a complex combination of issues centering on major engineering challenges.

1. Over-sampling wide bandwidth signals demands a high precision clock source. Thus a high performance clock and PLL combination is an essential element of the CT $\Delta\Sigma$ solution. It is only made possible by the development and integration of a novel inductive LC tank circuit. Low jitter clock sources are a costly component required in all high performance data conversion systems. Integrating this key network simplifies the application of the CT $\Delta\Sigma$ ADC, saving the design time and cost of specifying suitable external devices.

2. Deep sub-micron CMOS processes have been essential to making high-speed $\Delta\Sigma$ modulators a reality. Firstly they bring benefits in lowering the cost of implementing complex multi-stage digital filters. Additionally they support high clock rates allowing wide input signal bandwidths. Now, in die area terms, $\Delta\Sigma$ modulators challenge the economics of pipeline ADCs. Finally, with inherently low noise sensitivity, $\Delta\Sigma$ modulators offer enhanced integration opportunities that would not normally be considered in pipeline converters.

3. A high-speed multi-bit 3rd order $\Delta\Sigma$ modulator is needed. Xignal has demonstrated the highest speed $\Delta\Sigma$ modulator block available on a standard CMOS process. This is a significant innovation step. It requires a tunable, high gain, continuous time loop filter stage. The tuning is especially important given

the need for a general-purpose design to work over a range of sample rates. Ensuring stability of such a design is challenging. Xignal's design ensures unconditional stability.

4. Most previous $\Delta\Sigma$ modulator development work has centered on switched capacitor loop filter designs. Implementing continuous time systems requires access to high performance analog blocks including high precision 14-bit feedback DAC and high gain op-amp stages. Such complex analog circuit elements are a challenging area of development in CMOS.

6 Conclusions

The advanced CT $\Delta\Sigma$ ADC solution that has been discussed shows tremendous promise. Initial results are very encouraging and better than the FOMs achieved by contemporary pipeline ADC designs. Furthermore, the CT $\Delta\Sigma$ architecture has a clear advantage with its ability to scale with CMOS process developments and yield further increases in efficiency and speed. This is a major benefit, as pipeline design will become ever more challenging within the restrictive confines of the CMOS process roadmap with its sub-bandgap threshold requirements.

The quantization noise simulation result illustrated by Figure 4 hints at the unrealized dynamic range of Xignal's multi-bit $\Delta\Sigma$ modulator. There is some room for further dynamic range improvements. But attaining the full dynamic range will demand extremely careful design, paying particular attention to thermal noise, the dominant noise source in this design today.

The ongoing development at Xignal shows that the ADC core can be successfully combined with input signal path components to provide a high level of integration. The digital processing (filtering and decimation) provided also shows that in future it is possible for system designers to tailor transfer functions for a given application. Finally, the elimination of external anti-alias networks and the inclusion of a high performance PLL significantly ease the design of a high-resolution high-speed sampling system.

Though complex, the technology deployed within Xignal's CT $\Delta\Sigma$ ADC has been implemented in such a way to be transparent to the user, making this one of the most complete and easiest-to-use data acquisition systems available. Its implementation challenges some of the key assumptions made by designers developing high performance data conversion systems and offers a unique alternative to the pipeline conversion route.

Acknowledgement

For their not insignificant contributions to the preparation of this paper, the author wishes to acknowledge the assistance and support of the whole ADC development team at Xignal Technologies.

By Mark Holdaway,
Product Marketing Director, Xignal Technologies AG.

Xignal Technologies AG.

Leipziger Str. 16
D-82008 Unterhaching
Germany.

Tel: +49 (89) 32 22 72 -0.

Web site: <http://www.xignal.com>