How a Sigma-Delta ADC Works

James M. **Bryant**, Analog Devices

DESCRIPTION

Most descriptions of how a sigma-delta ADC works start with a few pages of partial differential equations and go downhill from there. This module eschews mathematics. It discusses **oversampling, noise-shaping, and digital filters and decimation**, and how these simple techniques enable the manufacture of high-resolution ADCs which are cheap, low-powered, linear and capable of self-calibration.

http://www.techonline.com/community/ed_resource/course/13409

A Σ–Δ ADC is simple but its mathematics is complex This lecture ignores the mathematics & concentrates on what actually happens

Sigma-Delta Analog-Digital Converters (Sigma-Delta ADCs) have been known for nearly thirty years, but only comparatively recently has the technology (high density digital VLSI) existed to manufacture them as inexpensive monolithic integrated circuits. They are now used in many applications where a low-cost low-bandwidth high-resolution ADC is required. There have been innumerable descriptions of the architecture and theory of Sigma-Delta ADCs, but most commence with a maze of integrals and deteriorate from there. In the Applications Department at Analog Devices we frequently encounter engineers who do not understand the theory of operation of Sigma-Delta ADCs and are convinced, from a study of a typical published article, that it is too complex easily to comprehend.
There is nothing particularly difficult to understand about Sigma-Delta ADCs, as long as you avoid mathematics like the plague. This lecture describes how they work in **non-math**ematical terms

A Σ–Δ ADC contains simple analog circuitry and complex digital circuitry.

It works by Oversampling, Noise Shaping and Digital Filtering (Decimation)

A basic Sigma-Delta ADC contains very simple analog electronics (a comparator, a switch, and one or more integrators and analog summing circuits) and quite complex digital computational circuitry. This digital circuitry consists of a digital signal processor (DSP) which acts as a filter (generally, but not invariably, a low pass filter [LPF]) and may perform other computations as well. But it is not necessary to know how the filter works to appreciate what it does.
To understand how a Sigma-Delta ADC works one should be familiar with the concepts of *over-sampling, noise shaping, digital filtering*, and, possibly*, decimation*

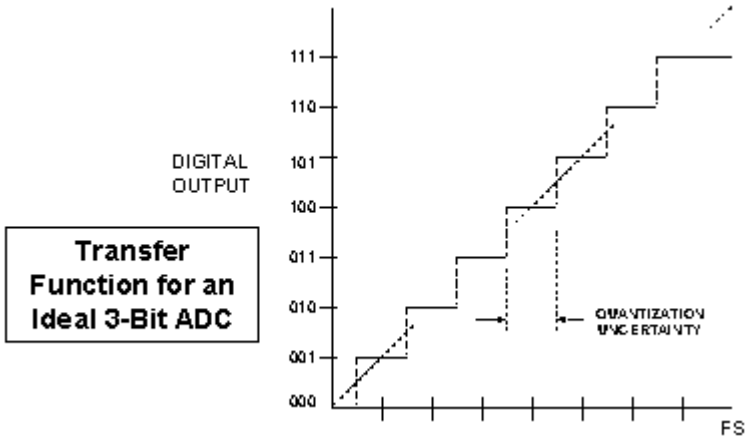

An ADC is a circuit whose digital output is proportional to the ratio of its analog input to its analog reference. The scaling factor between the analog reference and the analog signal is frequently unity and the digital output then represents the normalised ratio of the two.

The diagram shows the transfer characteristic of a 3-bit unipolar ADC.

The input to an ADC is analog and not quantised, but its output, of course, is quantised. The transfer characteristic therefore consists of a number of horizontal steps. When considering the offset, gain and linearity of an ADC we consider the line joining the midpoints of these steps.
Digital full scale (all 1s) corresponds to 1 lsb below the analog full scale (the reference or some multiple thereof). This is because the digital code represents the *normalised* ratio of the analog signal to the reference

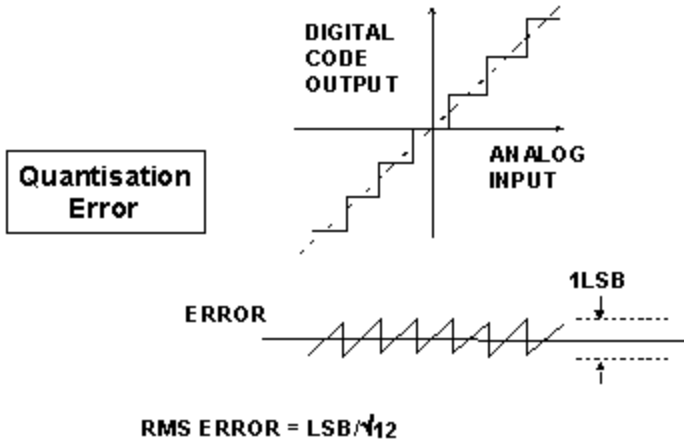

In an ideal ADC transitions take place at ½ lsb above zero and thereafter every lsb until 1½ lsb below analog full scale. Since the analog input to an ADC can take any value while the digital output is quantised there may be a difference of up to ½ lsb between the actual analog input and the exact value of the digital output.

This is known as the "quantisation error" or "quantisation uncertainty".

The word "error" might imply a deficiency. In fact the quantisation error is a fundamental part of the operation of an ADC, and has a mean value of the square root of one twelfth of an LSB

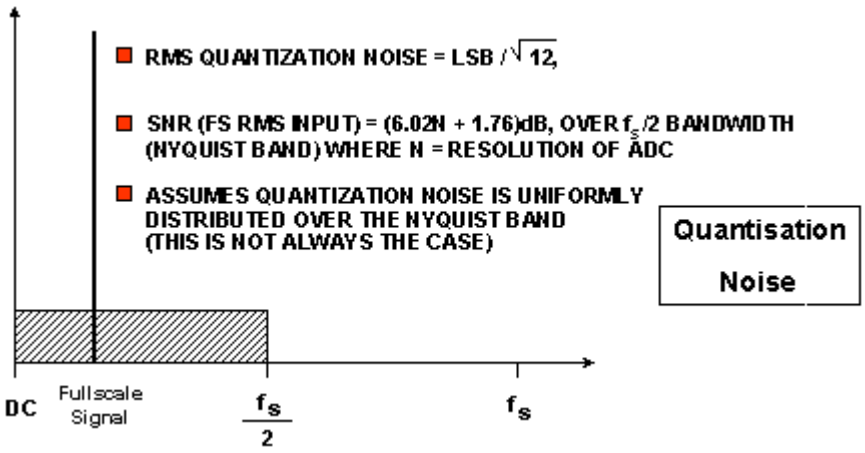

In sampled data applications, that is to say applications where the input to the ADC is an AC signal, this quantisation error gives rise to "quantisation noise".

For a perfect N-bit ADC the quantisation noise, relative to a full-scale sine input, is -(6.02 N + 1.76) dB.
In the diagram this noise is shown uniformly distributed over the band from DC to FS/2 (the “Nyquist band”).
As we shall see, this is not always the case

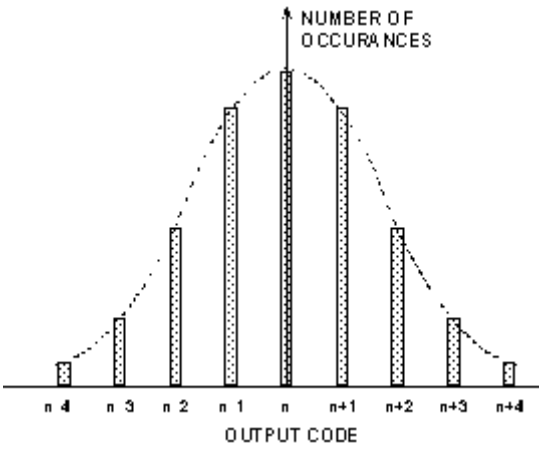

If we apply a fixed input to an ideal ADC we will always obtain the same output and the resolution will be limited by the quantisation error.

Suppose, however, that we add an AC signal to the fixed signal, take a large number of samples, and prepare a histogram of the results. We will then obtain something like the diagram. If we calculate the mean value of a large number of samples we will find that we can measure the fixed signal with greater resolution than that of the ADC that we are using. This procedure is known as over-sampling and the AC signal is known as a "dither" signal.

With the simple over-sampling described here the number of samples must be doubled for each bit of increase in resolution

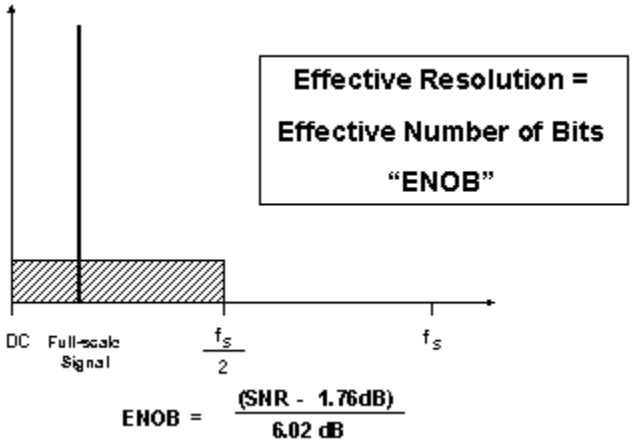

Let us consider the technique of over-sampling in the frequency domain.

A perfect classical N-bit sampling ADC has an r.m.s. quantisation noise of one upon root twelve of an LSB, uniformly distributed in the Nyquist band of DC - fs/2 (where fs is the sampling rate) and therefore its SNR with a full-scale sinewave input will be (6.02N + 1.76) dB. If the ADC is less than perfect and its noise is greater than its theoretical minimum quantisation noise, then its effective resolution will be less than N-bits.

Its actual resolution (often known as its Effective Number Of Bits or **ENOB**) will be given by:

(SNR-1.76dB) / 6.02 dB

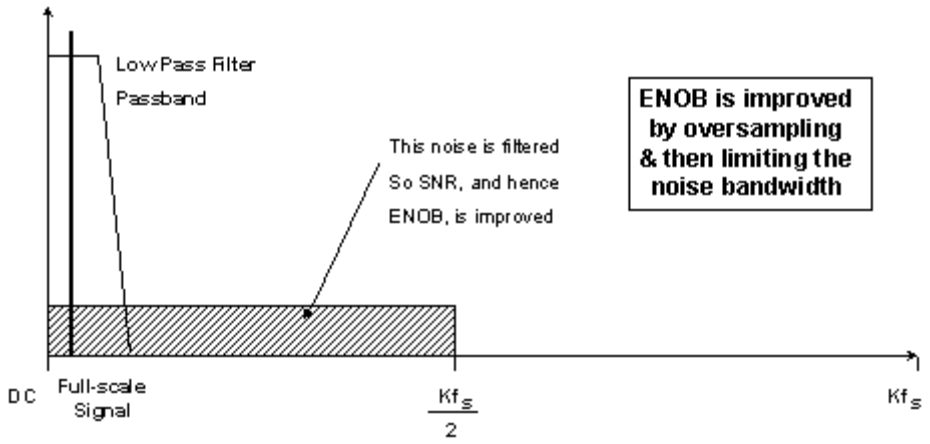

If we choose a much higher sampling rate (K times fs) the total quantisation noise is distributed over a wider bandwidth. By applying a low pass filter (LPF) to the digital output data stream we remove much of the quantisation noise but do not affect the wanted signal - so the ENOB is improved.
We have accomplished a high resolution A-D conversion with a low resolution ADC

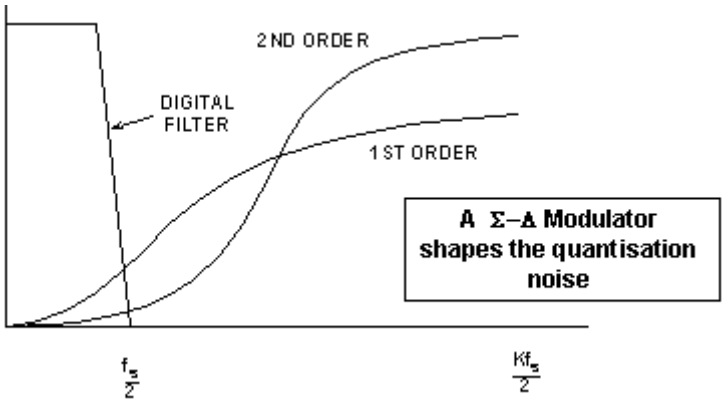

But, as we mentioned earlier, the sampling rate must double for each bit of increased resolution. If we simply use over-sampling to improve resolution we must over-sample by a factor of 2 ^ N to obtain an N-bit increase in resolution.
The Sigma-Delta converter does not need such a high over-sampling ratio because it not only limits the signal passband but also shapes the quantisation noise so that it falls outside this passband

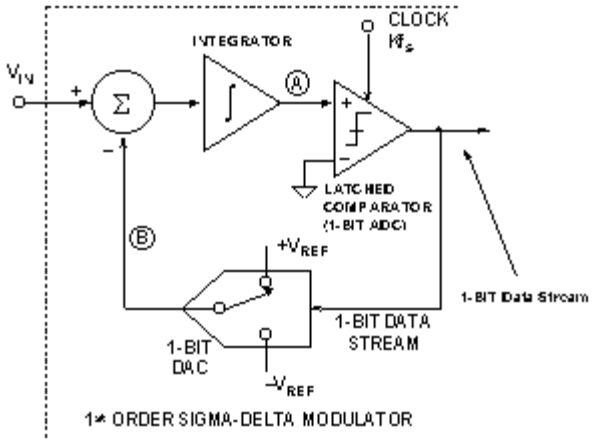

If we take a 1-bit ADC (generally known as a comparator), drive it with the output of an integrator, and feed the integrator with an input signal summed with the output of a 1-bit DAC (a switch) fed from the ADC output, we have a first-order Sigma-Delta modulator.
Its output is a data stream which has shaped quantisation noise - concentrated in the HF end of the Nyquist band

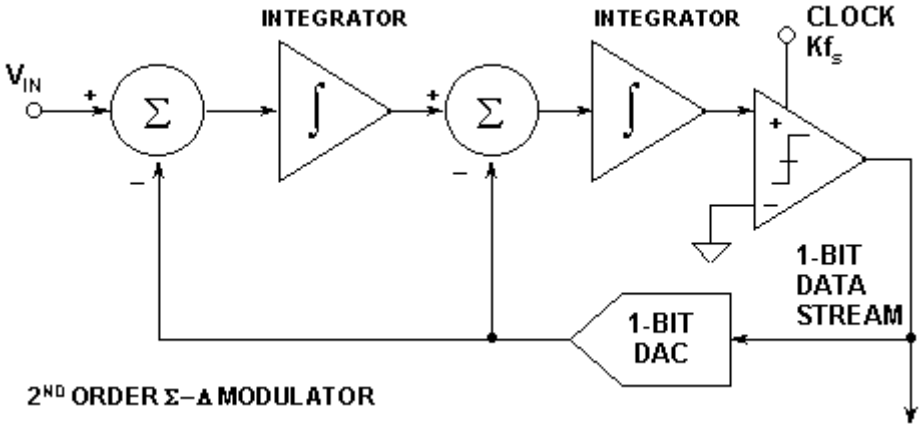

By using more than one integration and summing stage in the Sigma-Delta modulator we can achieve higher orders of quantisation noise shaping and even better ENOB for a given over-sampling ratio. A second-order Sigma-Delta modulator is shown in this diagram, but third-, fourth- and fifth-order modulators are also commonly used.
Third, and higher, order Sigma-Delta modulators were once thought be potentially unstable at some values of input - recent analyses using *finite* rather than infinite gains in the comparator have shown that this is not necessarily so for third-order Sigma-Delta modulators. But even if instability does start to occur it is not important, since the DSP in the digital filter can very easily be designed to recognize the onset of instability, and to respond to prevent it

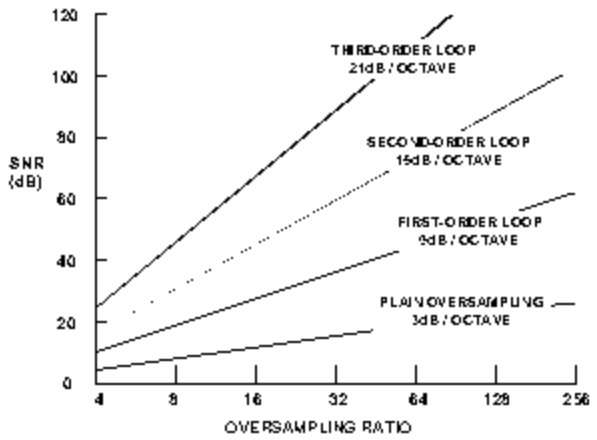

This diagram shows the relationship between over-sampling ratio and ENOB for first-, second-, and third-order Sigma-Delta modulators, and for simple over-sampling.
It is very evident that the use of a Sigma-Delta modulator gives greatly improved performance over systems using over-sampling without noise shaping

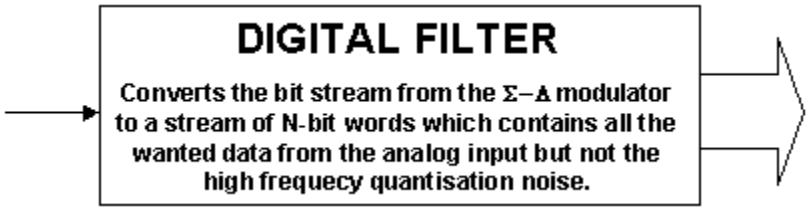

The output of the Sigma-Delta modulator is a stream of bits, 0s and 1s, synchronous with the sampling clock. This data stream contains all the data we require about the analog input to the ADC, but it also contains a lot of high frequency noise.
It is possible, indeed quite simple, to build a special purpose digital computer which processes such a data stream and provides as an output a series of multi-bit words which represent the input to the ADC. This computer is called a digital signal processor (DSP) and in the present case digitally performs the function of a low-pass filter (LPF).
How the digital filter operates is beyond the scope of this lecture. (It is one of the places where mathematics might get in!) However the notes for Analog Devices recent “Mixed Signal Seminar”, which are available from Analog Devices, give a more detailed overview of the operation of such filters

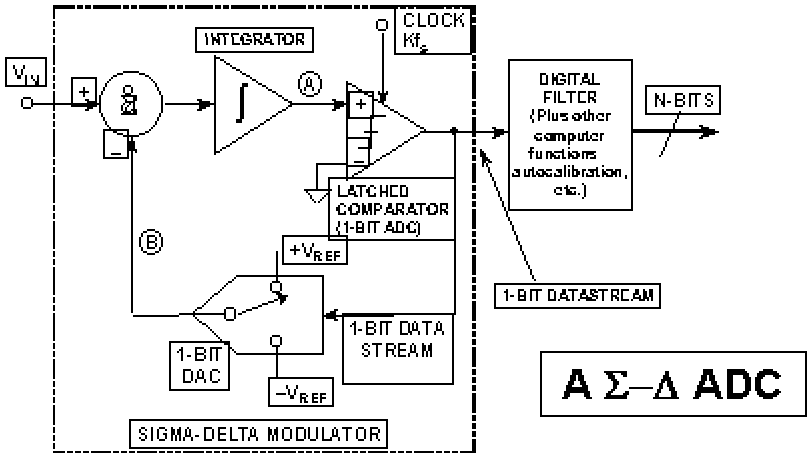

This combination of Sigma-Delta modulator and digital filter forms a Sigma-Delta ADC. The version illustrated uses a first-order modulator, but, as we have already said, modulators up to fifth-order may be used.
The digital filter is a special-purpose computer. It is frequently given other tasks besides its basic filter function. These may include ensuring the stability of high-order loops (as we have already mentioned) and autocalibration of the ADC

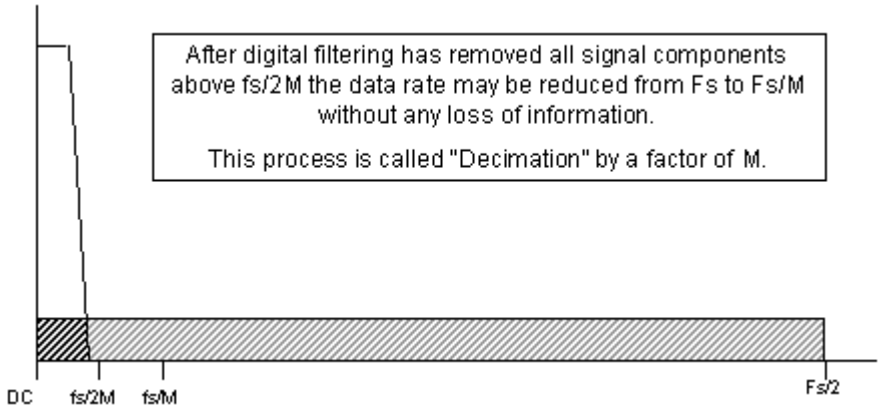

Since the bandwidth is reduced by the output filter, the output data rate does not need to be as high as the original sampling rate to continue to satisfy the Nyquist criterion.

In nearly all Sigma-Delta ADCs the output data rate is reduced to a fraction of the sampling rate by passing every Mth result to the output and discarding the remainder. The process is known as "decimation" by a factor of M. Despite the origins of the term [*decem* is Latin for ten], M can have any integral value, provided that the output data rate is more than twice the signal bandwidth.

Decimation results in no loss of information

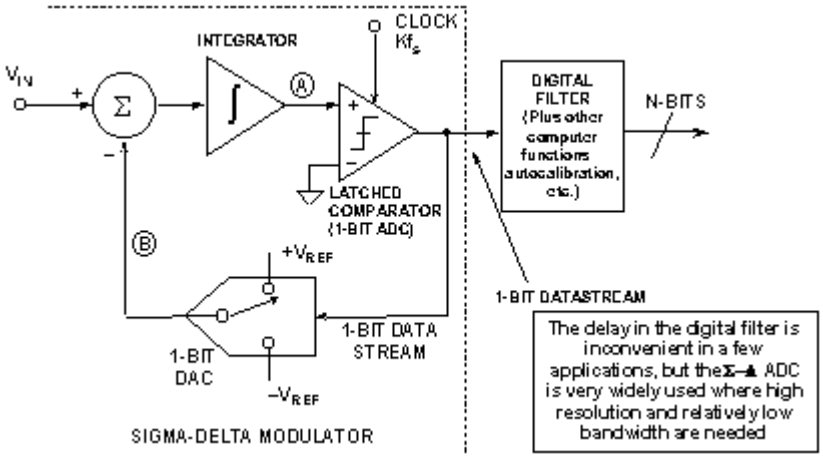

We have now seen that **although the mathematics of a Sigma-Delta A-to-D may be complex, its basic operation is easily understood**.
Its architecture is well-suited for high resolution and relatively low signal bandwidths – it is also fundamentally inexpensive.
But a digital filter has a fundamental **delay** in its characteristic – so data travels from the input to the output of a Sigma-Delta A-to-D relatively slowly compared to some other types of ADC and this is inconvenient in a few applications. Nevertheless Sigma-Delta A-to-Ds are very popular and widely used in many different applications.
*
The remainder of the lecture considers two useful variations on the common basic architecture

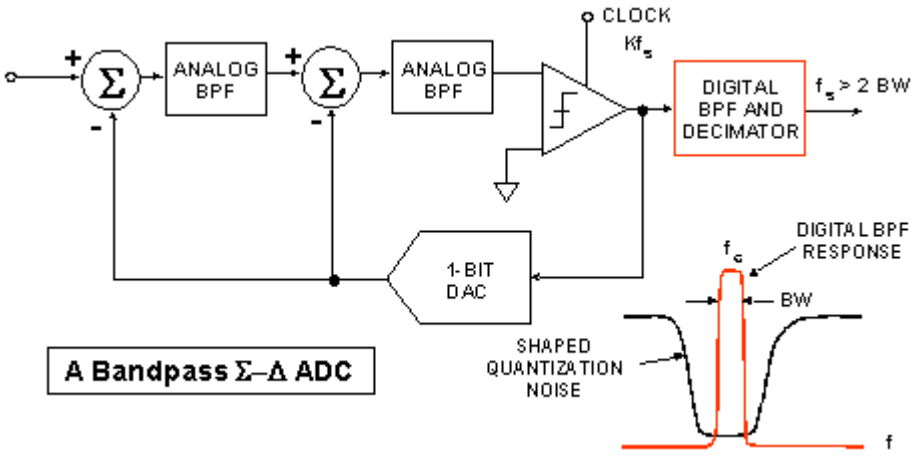

The Sigma-Delta ADCs that we have described so far contain integrators, which are low pass filters, so they have a passband extending from DC, and their quantisation noise is pushed up in frequency. At present all commercially available Sigma-Delta A-to-Ds are of this type (although some which are intended for use in audio or telecommunications applications contain bandpass rather than lowpass digital filters to eliminate any system DC offsets).
But there is no particular reason why the filters of the Sigma-Delta modulator should be low pass, except that traditionally A-to-Ds have been thought of as being baseband devices, and that integrators are somewhat easier to construct than bandpass filters. If we replace the integrators in a Sigma-Delta A-to-D with bandpass filters (BPFs) the quantisation noise is moved up and down in frequency to leave a virtually noise-free region in the pass-band. If the digital filter is then programmed to have its pass-band in this region we have a Sigma-Delta A-to-D with a **bandpass**, rather than a low pass characteristic.

Although they are not yet manufactured as general-purpose devices, such A-to-Ds are ideally suited for use as converters in ASICs for digital radio receivers, medical ultrasound and a number of other applications

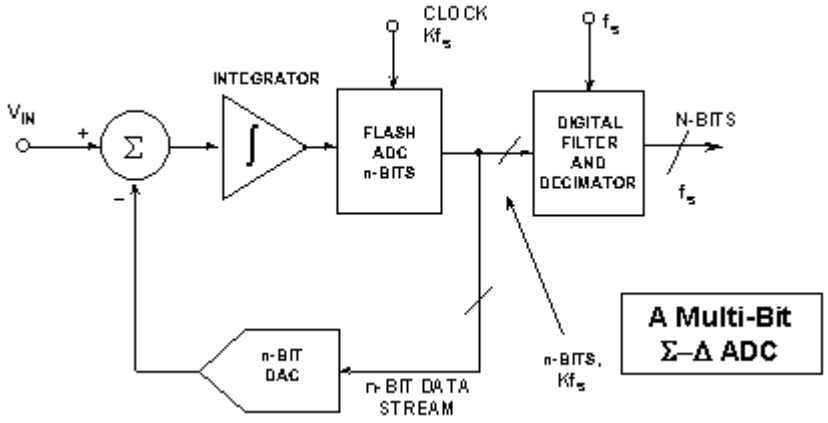

There are also no particular reasons, except simplicity and intrinsic linearity, why the comparator and switch of a Sigma-Delta A-to-D should not be replaced with a more complex **multi-bit** ADC and DAC.
This architecture allows greater bandwidth for a given resolution and sampling frequency, at the cost of greater analog complexity and possible non-linearity.
There are many techniques for minimising this non-linearity and such A-to-Ds are manufactured and quite widely used

Σ–Δ ADCs are
High Resolution,
Limited Bandwidth,
Intrinsically Linear,
Inexpensive &
***Easily Understood* !**

A Sigma-Delta A-to-D works by **over-sampling**, by using simple analog filters in the Sigma-Delta modulator to **shape quantisation noise** so that the SNR in the bandwidth of interest is much lower than would otherwise be the case, and by using high performance **digital filters and decimation** to eliminate noise outside the required passband.
Because the analog circuitry is so simple and undemanding it may often be built with the same digital VLSI process that is used to fabricate the DSP circuitry of the digital filter, and because the basic A-to-D is usually **1-bit** (a comparator) the technique is inherently **linear**.
Although the detailed analysis of Sigma-Delta A-to-Ds involves quite complex mathematics their basic design can be understood without the necessity of any mathematics at all.

James M. **Bryant**
Head of European Applications
Analog Devices Inc.
January 2001
James.Bryant@analog.com