22bit Sigma-Delta ADC Theory and Circuit Operation

J. G. Pett

Introduction

A basic Sigma-Delta ADC is shown in Fig.1 and has an action essentially the same as a Charge-Balance type. This is termed a 1st order modulator. Additional integrators can be added between the first integrator and the comparator, where Fig.2 shows a 3rd order modulator. The following sections explain firstly the basic theory of operation of a SD type of ADC and then give a detailed description of the modulator circuit. The digital filter is described in the final section.

Theoretical Concepts

Oversampling

Any ADC can be described as an electronic circuit that quantises a given input signal. This means that the digital output has a number of output codes, each of which describes a particular interval of the input signal. For a perfect ADC of n bits, each interval will be equal to $V_{fsr}/ 2^n$ [V_{fsr} = the full-scale range of the input signal]. If a perfect transition from one output code to the next occurs as the input signal increases, then it becomes possible to interpolate between each interval using averaging techniques and hence improve the resolution beyond that of the n bit ADC. Such methods are termed oversampling. [In practise, the input signal must frequently traverse at least two adjacent intervals to ensure correct averaging. This is assured for most cases by the signal noise, the ADC internal noise or an externally applied "dither" signal, which sweeps the input over a narrow range of bit intervals, or transitions.]

Normal averaging of 10 output values allows each interval to be divided by 10, effectively increasing the resolution by more than 3bits, but this requires sampling the signal 10 times faster to retain the same signal bandwidth. Alternatively, we can express this as reducing the quantising "noise", or uncertainty. Averaging 100 values will improve the noise by 40dB and so on. Oversampling by 500 will reduce the quantisation noise by 54dB, however this also means that ADC bandwidth is traded against improved resolution.

Consider an example from the 22bit ADC. For a 1kHz bandwidth signal to digitise, we must sample at 2kHz [Nyquist criteria]. If we now oversample at 1MHz to improve the resolution, we will obtain approx. 9 extra bits of resolution and improve the signal to noise ratio by 54dB. [Note : the definition of oversampling ratio is $f_s/f_{Nyquist}$.]

Noise Shaping

Provided that the input signal is ***not*** a stationary [DC] value, the above quantisation noise can be described as 'white noise' with a bandwidth of $f_s/2$.

The basic Sigma-Delta modulator is a form of charge-balance ADC having a 1bit quantiser [the comparator] and a 1bit very-high-precision DAC. The quantisation noise is that coming from the 1bit DAC and is large [N/S is approximately 0dB]. This noise is integrated by the

first integrator and hence modifies the theoretical white noise spectrum, reducing it where the integrator has internal gain i.e. below the gain = 1 point. This process is termed noise shaping. In the 22bit case, if the integrator gain = 1 point is set at 25kHz then the noise reduction is about 20dB/decade and at 1kHz is approx. = 28dB. Adding this to the above 54dB gives an overall N/S ratio of about -82dB. Therefore, with a single integrator filtering the quantisation noise of a very simple 500x oversampling, 1bit ADC we can obtain a resolution at 1kHz of about 1part in 12,000 or nearly 14bits. Extra stages filtering the quantisation noise are possible and can improve the resolution by about 28dB for each additional integrator. Again using the 22bit case we obtain 54 + (3x28) = 138dB or >22bits.

Noise shaping using filters in the forward path of the charge-balance circuit, along with a high oversampling ratio, constitute the basic theoretical concepts on which Sigma-Delta ADCs work.

Filtering

It remains to filter the wanted output data from the 1bit stream of data at 1MHz. This is done using a low-pass digital filter, which essentially removes the quantisation noise above the wanted band [DC to 1kHZ] while increasing the output word length. In order to suppress noise above the signal pass-band, a low pass filter having approximately the same attenuation rate as the noise shaping filter is required, to match the rising noise-shaped spectrum. This ensures that all signals above the filter cut-off frequency are kept below the base-band noise floor.

If the input signal spectrum contains components above f_s /2, aliasing will occur. Therefore an analog input filter must lower any such signals to below the ADC noise threshold. For traditional sample rates just above the Nyquist rate such a filter becomes a major design limitation. However, the SD ADC actually samples at say 500x Nyquist, which means that adequate attenuation can be achieved with simple passive single or two pole filters. This is particularly important for very high performance ADCs since input signal quality can be maintained.

Functional Description of 22bit SD Circuit Diagram

Isolation and EMC

The whole of the modulator circuit is floating with respect to ground and is contained within a guard box. This method assures a common-mode rejection ratio of greater than 150dB at 50Hz. To provide adequate EMI suppression, a capacitive coupling network is placed between the input cable screen and ground, followed by two isolating hf chokes at the single-ended buffer input. This protection has been shown to provide >2.5KV burst test immunity.

Input Buffer and Anti-Alias Filter

The high input impedance buffer amplifier IC16 is preceeded by a two-pole anti-alias filter [break freq. = 5kHz] giving > 94dB attenuation at the sampling frequency of 1MHz. Any input signals at 1MHz +/-1kHz should therefore be < 100mV [-40dB of 10V] to avoid aliasing, and this will normally be easy to ensure.

Integrators

The 1st summing integrator is IC10, which combines the input signal via a 70K and the 20K resistors switched by IC8, which together constitute the 1bit precision DAC. Note that all these resistors are Vishay precision types. In order to circumvent the recently discovered offset instability of IC10 [Ti/BB 627], an additional chopper-stabilised amplifier ICXX [LTC 1250] has been applied, such that the combined amplifier pair now has both excellent DC offset stability [<0.05ppm/degC] and high frequency performance. An additional and important bonus is that the 1/f noise of the 627 has been suppressed, leading to improved long-term stability.

The 2nd and 3rd integrator stages are constituted by IC15 A+B which are both configured as ***non-inverting*** integrators. This design has two advantages. The first is that each circuit acts as a unity-gain stopped-integrator above the gain 1 break frequency and the second is that if saturation occurs due to excessive input signal, clamping can be achieved, [using analog switches IC14], to give unity gain from input to output of the pair. Under such clamped conditions the 3rd order configuration [basically unstable] reverts to the 'always stable' 1st order form and hence continues to function "normally" until the overload is removed. Note that the overload detection and clamp drive is described later.

Comparator and Dither Generator

The output of the 3rd integrator drives the comparator IC12, the output of which is clocked at 1MHz into IC7. [Note that IC7 decides the polarity of the 1bit DAC, thus closing the self-oscillating loop.]

A triangular 960kHz "dither" signal is injected into the + ve input terminal of the comparator for "idle-tone" suppression. The frequency generator is IC11 and P2 adjusts the output to 960 +/- 5 kHz. IC13A buffers this signal. P3 is set such that the pk-pk amplitude of the signal at TP10 is approx. doubled. This effectively randomises the idle-tone spectrum.

Note - An 'idle-tone' is a part of the modulator output noise occurring at a particular frequency, which falls within the required pass-band. Often a number of such tones can occur together. They are particularly noticeable with DC inputs. They can exceed the pass-band noise floor by up to 20dB and hence become problematic. Idle-tones decrease with increasing order of the modulator. The above method of suppression has never been reported in earlier literature.

Precision References

The basic precision-voltage-reference design is a modified "Spreadbury" circuit, using IC5 and IC6. IC6C provides internal temperature stabilisation of the LTZ1000A buried zener, where R32 sets this temperature to approx. 35degC. Note that R33 is an initial set-up component = 3.3K which allows the LTZ 1000A to 'burn in' for 15days at 105degC. After this period R33 should be removed for normal operation. This 'burn-in' removes zener drift, which would otherwise occur during the first year of operation. IC6D provides the feedback to T2, which sources the zener. Note that the basic "Spreadbury" circuit is very sensitive to EMI and that extra ceramic decoupling capacitors have been added to suppress these effects. The 7.1V output at TP4 connects to an attenuator R70+R75 [2x20K] where the centre point of the 20K pair is held at zero volts by IC6B. IC6A buffers approx. +5V to power IC8, thus providing the precision positive voltage for the 1bit DAC. IC6B provides a permanent -5V to the

combination R73, R74B. Note again the use of L3-5 and extra EMI suppression capacitors. The output impedance at high frequency of the basic +5V reference is far too high and heavy decoupling is necessary to prevent signal droop occurring during "shoot-through" operation of the CMOS output pairs of IC8. Fortunately, the multi-layer ceramic capacitor C20, 2.2uF solves this limitation. P4 provides the only fine adjustment of DAC output symmetry but as such it is designed to incorporate all other circuit asymmetries and is ***not*** intended to ***balance*** the actual +/-5V reference outputs.

Clock Source, Logic and 1bit DAC Switching

The basic clock source is derived from a 16MHz Xtal controlled oscillator Q1, which has heavy supply-line decoupling to minimise clock "pulling" and phase jitter. This Xtal type must not be replaced, as it is critical to obtaining the overall performance. The output of the oscillator is buffered by IC2B and inverted by IC2A, thus allowing both edges of the signal to be used. The 16MHz clock feeds an FPGA, IC3, which is programmed to divide this input to provide the 1MHz clock for IC7 [pin 28]. At the same time a 62.5nsec 'guard pulse' is generated [pin 9 of IC7B], overlapping the 1MHz edge, which serves to clamp the 1bit DAC output to zero during basic clock transitions, and also ensuring precisely equal switching surfaces in the 1bit DAC output irrespective of pulse sequencing.

IC7A latches the comparator value with the 1MHz clock and returns this value to the FPGA [pin 31]. If the value is a +1 then a 250nsec pulse is output [pin 22] to the fibre Tx. A -1 outputs a 750nsec pulse. This allows the leading edge of these pulses to be use remotely for re-generation of the 1MHz clock and a sampling of the state of the pulse at 500nsec allows either a +1 or a -1 to be detected.

The FPGA detects if the modulator is in saturation by detecting more than 15 continuous output pulses of the same sign. Should this occur, the FPGA outputs a 1µsec pulse [pin34] to short circuit the integrator capacitors of the 2nd and 3rd integrator stages, thus "resetting" them.

IC9 combines the 'direction' information from the latched comparator signal [IC7A] with the 62.5nsec 'guard pulse' to drive the parallel connected XOR gates. Each pair of gates functions as a precise analog switch connecting either the positive or the negative 5V reference voltage to R73 + R74B and thus injecting a precise charge quanta into the 1st integrator capacitor.

Precision Switches

IC8 B+C acts as a non-inverting buffer, where the totem-pole output stage has almost symmetrical and low output impedance [$R_{on\text{-}up}$ = 11ohms, $R_{on\text{-}down}$ = 7ohms]. IC8A+D acts as an inverting buffer. [During the early development, this pair was powered between 0V for V+ and -5 V reference for ground. Initially it may be easier to understand the circuit operation in such a connection]. This second pair connected as shown, with R74B pulling down to -5V reference, actually achieves the same result, but now all logic is referred to ground. The overall effect of this somewhat complicated switch is that during the 'guard pulse' IC8B+C output is closed to ground and IC8A+D output is closed to +5V reference, which cancels the current sourced from -5V reference via R74B. When a positive current must be injected into the integrator input the output of IC8B+C is at +5V reference. However, when a negative current must be injected IC8A+D must close to ground thus allowing the current to be sourced directly via R74B. The different 'up' and 'down' switch resistances, the precision resistor tolerances, as well as the two reference voltages can cause slight imbalance between

the currents injected into the integrator. Series trim resistorsR57+58+62 allow an overall symmetry to be achieved [note that R62 is set to 18ohms during initial assembly]. Fine adjustment during calibration, plus adjustment of the REF BAL potentiometer allows precise zeroing to be achieved.

Great care was taken to design a minimal overall delay for each switch circuit with very symmetrical and short rise and fall times [IC's 8 +9]. Hence, the logic family and the pcb layout should be considered vital to the overall accuracy and changed only after extensive measurements. The TC of this precision switch dominates the temperature coefficient (TC) of the whole ADC. Hence, a Peltier device has been incorporated to control the overall temperature of this critical area to <1degC.

Temperature Control

The temperature control circuit aims to stabilise the aluminium block containing IC5 and which is thermally connected to ICs 7,8 and 9 [the critical switching components]. The block also contains the temperature sensor IC4, which must be tightly coupled thermally to this block. This sensor has a temperature coefficient (TC) of 5mV/degK and also outputs a reference voltage of 2.5V. To set a given temperature, P1 [SET TEMP] takes a fraction of this reference voltage, which is then buffered by IC1C and can be observed at TP3. For 25degC, this voltage should be set to 1.49V. The actual temperature, from IC4 pin5, is compared with this set value by IC1D. This difference is also amplified, such that at TP2 the error signal can be observed with a sensitivity of 1V/degC.

IC1B acts as the loop compensation amplifier, providing unity proportional gain and an integration term starting at 5mHz. A differential term is provided by R9, C69 [33mHz], which is stopped by R10 at 330mHz. This stabilises the overall control loop where the basic system time constant (i.e. Peltier+heatsink to sensor, via the aluminium block), is about 100seconds [1.6mHz]. Finally, to avoid excessive overshoot, a limiter circuit, consisting of R8 + D2,D3 clamps the excursion at IC1B pin7. The output of this 3-term controller feeds an output power amplifier connected to the Peltier element and a series resistor [R7 on the power supply card]. IC1A acts, open loop, on a pair of emitter follower FETS [T1, T2 on the power supply card], which then pass current through the Peltier element, thus either heating or cooling the surface connected to the aluminium block. Note that the heating and cooling time-constants are only approximately equal. The current measured across the series resistor R7 is fed back to the input of IC1A thus closing the local feedback. While somewhat crude, this simple circuit ensures no shoot-through in the output pair and provide a reasonable current control feed to the Peltier element. Note that voltage control of this element would not take into account the rise in internal resistance due to self-heating. Since this output circuit is a current source, no ripple filtering or series regulator is needed on the +/-4V supply lines to the Peltier.

The overall control of block temperature, which can be observed at TP3, gives fast slewing to temperature [typically less than 3 minutes] and a final stability well within 1/100th of a degC.

Power Supplies and Fibre-Optic Driver Card

The power supplies to the SD card are very conventional and employ standard IC regulators, where the +5V is derived from the +15V supply. This ensures correct start-up [anti-latchup] for the CMOS ICs. Test points and LEDs are provided but the circuit merits no further description. The 'mains' transformer however is a special low field design, with twin electrostatic screens. The inner encloses the 230V primary, which together with the input

mains filter, limits 'mains generated' EMI, while all the secondaries are enclosed by the second screen. This screen must be connected to the metal guard box uniquely [note the earlier drawings must be modified to reflect this!!].

The fibre-optic driver, which consumes ~50mA, is mounted on this card to avoid having this dissipation inside the temperature-controlled guard box. It consists of Schmitt-trigger IC3, an o/c driver and the output LED IC1 which drives the plastic [50m] or silica [500m] cable.

Digital Filters

A number of digital filters have been designed to extract the data from the above modulator and input this data to some type of computer.

The basic function required is to attenuate all frequencies, contained in the 1MHz, 1bit data stream, above 1kHz [with about -60dB/dec roll-off]. At the same time, the data needs to be averaged such that the number of significant bits increases to beyond the noise floor. In the 22bit case this means about 24bits.

A number of digital filter types have the potential to provide this performance but some trade-offs occur with each type. A FIR type can provide the attenuation rate quite easily but often incurs a long 'group delay' [i.e. the fixed time for the data to pass through the filter and which can cause problems in real-time usage]. The sinc^n filter employed is simple in hardware terms and is among the fastest.

The basic design, employed in all the filter types designed to date, is a cascaded moving average filter. Mathematically, a moving average filter is described by the sinx/x function or sinc function. Cascading 4 such filters [sinc^4], provides (4 x -13.3)dB/dec which approximately matches the noise-shaping characteristic of a 3rd order modulator. The output data rate of such a filter can be at the input rate, namely 1MHz but clearly no useful information is contained at this rate, since the cut-off frequency is at 1kHz. It is normal to decimate (or discard non-useful data) the output data rate down to a more useful value. In the 22bit case, a 5kHz maximum rate was chosen. A number of design details for the particular computer interface must be taken into account and appropriate measures taken to ensure data integrity. All these design details are fully described in the Filter Design documents [Ph. Semenaz]. This design allows an external trigger source to synchronise the data acquisition, which is a feature not often available with SD ADCs.

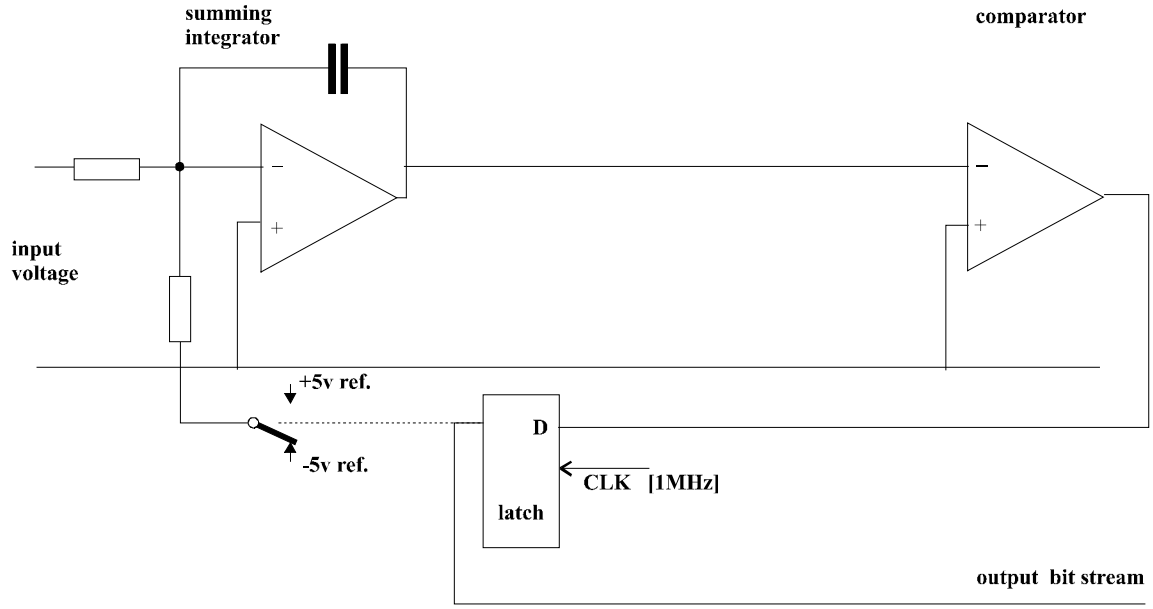

Figure 1. 1st order Sigma-Delta Modulator

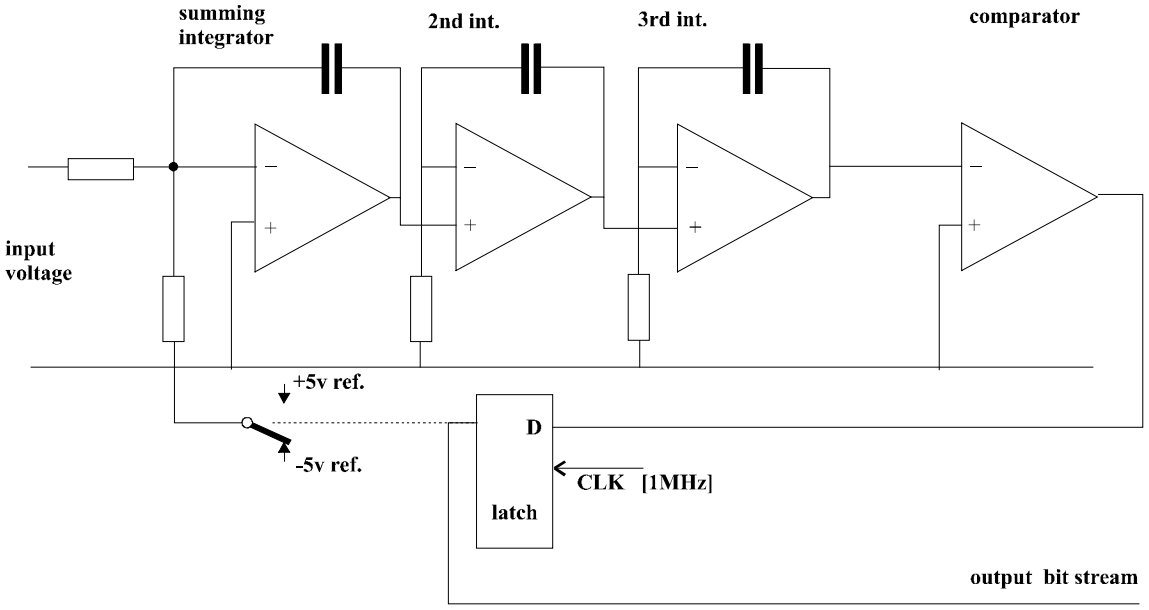

Figure 2. 3rd order Sigma-Delta Modulator