Audio Engineering Society

Convention Paper

Presented at the 114th Convention
2003 March 22–25 Amsterdam, The Netherlands

Design techniques for high-performance discrete A/D converters

Bruno Putzeys[1]

[1]Philips Digital System Labs, BE3001 Leuven, Belgium
bruno.putzeys@philips.com

ABSTRACT

An old 1-bit D/A converter idea using voltage switching is reintroduced. Some shortcomings of the original design are discussed. A new design is proposed that addresses these issues. A one-bit noise shaping ADC with DSD specs is built around this circuit serving as the feedback D/A. Theoretical noise performance and real performance are shown and compared. A roadmap for improving the present design is outlined. It is concluded that this conversion technique is much more promising than commonly thought.

1. INTRODUCTION

The present draws on work done, in parallel and independently by Green et al.[1] and by the author in the early 1990's on the construction of a 1-bit audio D/A converter using voltage pulse switching. The result of Green's work became available as a demonstration board from Crystal Semiconductor and provided, for the time, groundbreaking sonic performance.
This "direct" technique has since fallen to the background, yea mostly dismissed, on account of the following two main reasons:

- the perceived excessive clock jitter requirements together with the significantly better jitter-tolerance of charge-transfer converters and
- the high level of integration possible with charge-transfer converters, including the reconstruction filter.

It was dissatisfaction with the sonic performance of even the best examples of these that led to the reopening of the case.

2. ORIGINAL D/A CONFIGURATION

As D/A converters a dual ACMOS flip-flop was used, the 74AC11074. The bit rate as produced by the CS4303 chip (the actual product) is fr=64•fs or some 3.072MHz.

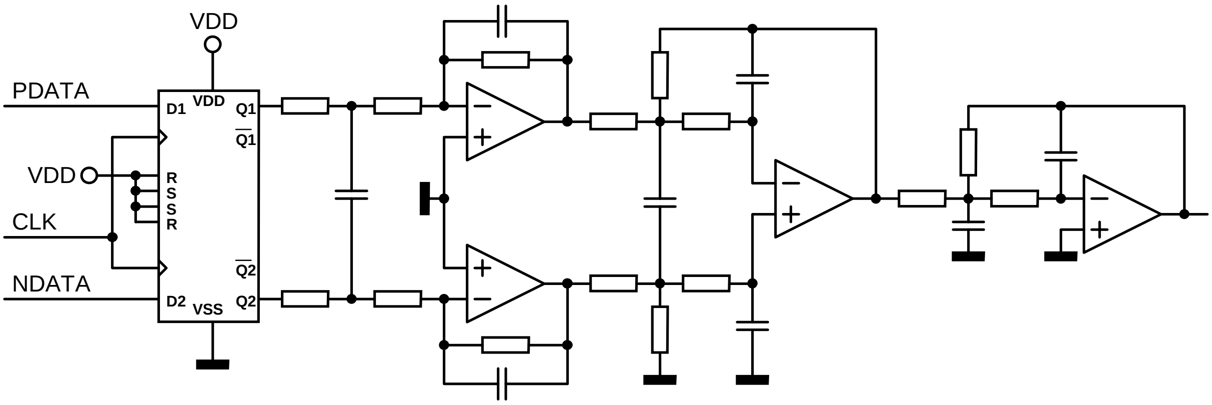

The filter is a 5+1 pole (first 2 real poles, then 2 complex pairs) active design executed differentially up to about halfway, where the first complex pair also performs the conversion to single-ended.

Positive and negative data inputs, as provided by the digital processor, are complementary versions of the 1-bit modulator output, both interspersed with “0”s. In doing so, each pulse period starts on a “0” regardless of the previous state. This insures that the pulse energy of each data period is independent of the previous state.

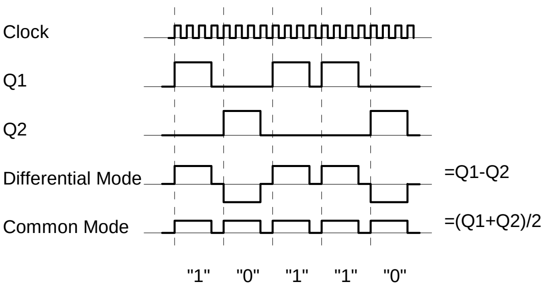

The differential mode signal, as considered by the filter, can be seen to have one “+1” state, one “-1” state and one “neutral” position which is assumed for one quarter of the bit time.

Unfortunately, the common mode signal is not zero. While it is true that it doesn’t contain any information, it is a rather hefty HF signal. This signal is supplied unattenuated to the first pair of op-amps which are required to attenuate it without exhibiting significant nonlinearities to affect the wanted signal. Only the best jfet op-amp will do here (the OPA627). Another high-quality op-amp (MAX427) was found to produce spectacular amounts of noise and distortion (0.1%). It stands to reason then that the OPA627 isn’t completely free from such problems either but that they are just “less obvious”. Indeed, whereas the CS4303 deltasigma modulator itself has an inband signal-to-noise ratio of 107dB, the evaluation board clocks in at only 100dB (105dB A-weighted).

3. MODIFIED D/A CONVERTER CIRCUIT

The circuit was modified to take advantage of the fact that after the dual flip-flop, not only the complementary return-to-zero signals are available, but also complementary return-to-one equivalents. Including these in the sum produces a differential signal of which both ends have a center detent.

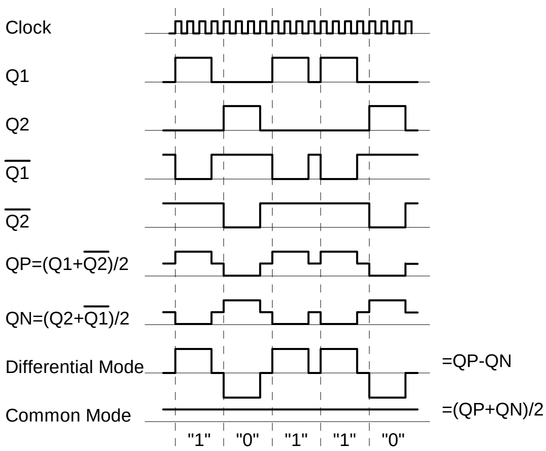

The common-mode component is now clean, save for any spikes caused by unequal rise and fall times. These spikes are so short (<1ns) that supplementing the first filter capacitor with a pair of much smaller ground-referenced ones will rid the signal completely of common-mode.

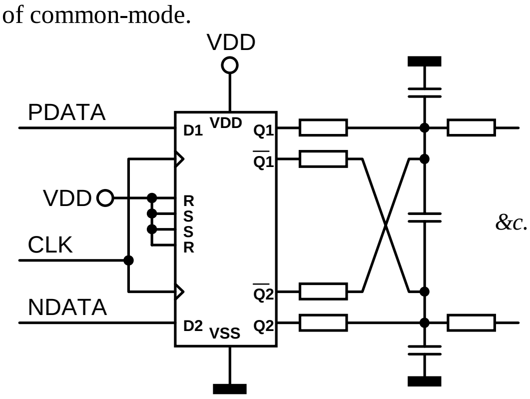

The main interest for this project, however, was not D/A converters. Selling one good D/A converter makes one customer happy. Selling one good A/D converter, however, makes many people happy by producing better-sounding recorded music. Apart from that, an A/D converter provides a better platform for testing the maximum performance capability of a given converter topology because the output (on which the actual measurement is done) is digital.

4. TOPOLOGY OF AN A/D CONVERTER

For the sake of simplicity, and because at the time DSD recording chains seemed a reality, the prototype was to use a 2.8224MHz 1-bit modulator with a noise-free bandwidth of 20kHz. A 6th order control loop is sufficient for this.

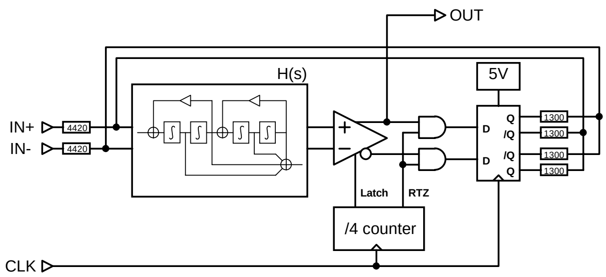

The odd resistor values derive from a desire not to leave one’s chair while stuffing a long-awaited circuit board. The comparator is equipped with a latch control which is released for a short time while the outputs are blanked (RTZ period). The timing allows the comparator to settle before the outputs are enabled again. With a delay of one clock period (one quarter of a bit time), these RTZ signals are reproduced by the “D/A converter”. This delay, and the partial zero-order hold function effected by the bit time are of no consequence if their presence is taken into account while synthesizing the loop function.
For reasons of noise gain, the input resistors are made large.

5. LOOP FILTER

The loop function consists of 3 pairs of poles, which can be optionally shifted away from DC.
Each pole pair consists of two noninverting integrators executed in a differential fashion.

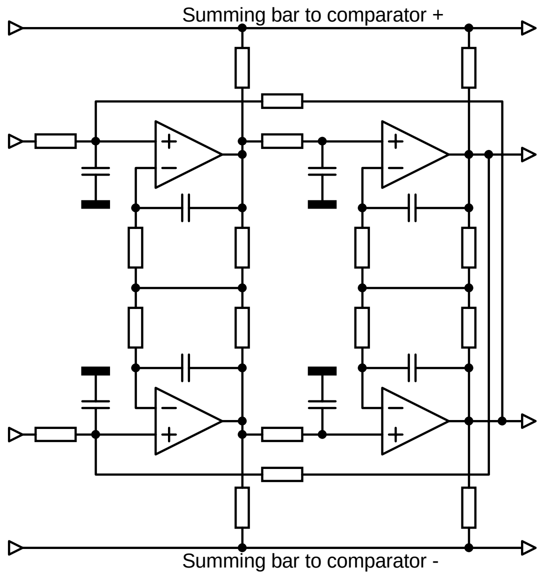

The very first op-amp in the chain should be selected both for excellent noise performance and large bandwidth (the group delay should be low compared to that of the digital chain). As noise shaper theory –and practice for that matter– has it, subsequent op-amps are much less critical and practically any run-of-the-mill type will do.
It will be obvious that when solving the desired loop function for the component values there are more unknowns than there are equations (even including the impedance scale factor, which should be chosen such as to obtain reasonable component values). The “missing equations” are solved through simulation or practice – that is, the integrator gains should be set such that under full-scale, maximum input frequency condition all op-amps are comfortably, but not too comfortably, within clipping. This insures stable startup and recovery from overload.
Also, the input resistors to the first integrator section are of course the combined DAC output resistors and the signal input resistors.
As for the synthesis of the loop function: there are countless methods known for this type of design. The simplest will consist of some basic maths hinging on the concept of a linearised quantizer. The most sophisticated will consist of several pages of highly involved maths culminating in a complex equation which is only solvable when the quantiser is approximated as a linear gain block. The difference between the two loop functions thus found is approximately 1.5dB. The difference in noise gain at the turnover point between any of both and the working noise shaper can be up to 10dB. Therefore, one should not have any moral qualms about opting for the simple method.

6. VOLTAGE REFERENCE

The voltage reference of the D/A converter is obviously the power supply rail it runs on. This supply should be as clean and as noise-free as practically possible. The noise performance of a simple 7805 style regulator will “only” deliver up to 112dB SNR. Luckily, it is not too difficult to improve on this. A filtered 5V reference, a fast low-noise op-amp, a transistor and good decoupling is all it takes. The electrolytic cap serves to damp the resonant circuit created by the output impedance of the op-amp/transistor combo (which increases with frequency and hence is inductive) and the ceramic decoupling cap. Due to the <90º phase margin of the active block the Q of this resonance is likely to be over infinity, so the circuit becomes unstable. The ESR of the elcap decreases the loop gain uniformly over a wide frequency range and can damp any resonance completely, yielding a perfectly stable circuit with excellent transient behaviour.

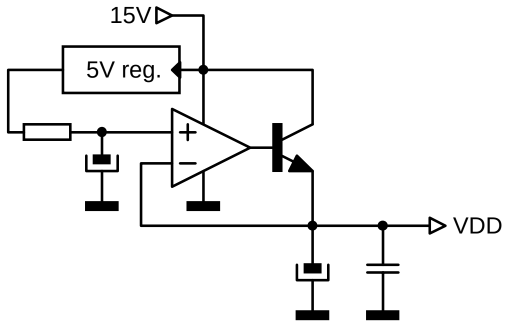

7. A MORE PRACTICAL DUAL FLIP-FLOP

When the 74AC11xxx logic series were introduced the two main reasons for this were the symmetrical “flow-through” style pinout and the decreased distance between the power supply pins.

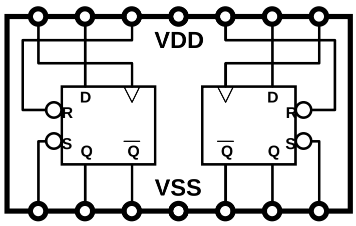

It was found however, that this does not provide enough of an argument to select this somewhat exotic single-supplier type.

Firstly, the pinout of the standard 74AC74 type is equally “flow-through” and symmetrical when read from top to bottom. In fact, the die inside the package is the same, rotated 90º.

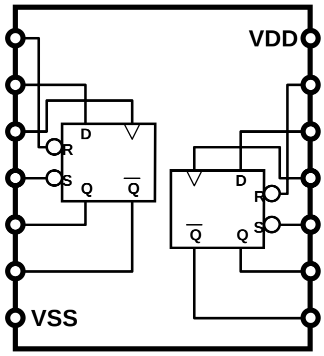

Secondly, the reduced pin distance between VDD and VSS does not bear out a practical advantage, for the inductance is mainly in the very thin bond wires. If a new pinout were to provide a lower loop impedance in the power supply circuit, the power pins would have to be located right next to eachother so that a portion of the bond wire inductance is mutual. The tiny inductance increase in the lead frame and in the PCB tracks can be made negligible compared to the bond wire inductance.

8. NOISE CALCULATION

Two op-amps were tried as first-in-chain, the OPA627 and the AD797. The OPA627 has an input referred noise voltage of 4.7nV/ √Hz, roughly equal to a 1.1kΩ resistor. Including the 10Ω output resistance of the flipflops, each feedback leg has an equivalent resistance of 570Ω. There are two sets of each, adding up to 3340Ω or 7.8nV/√Hz.

The noise gain is 1.15 times, making for an output referred noise of 8.9nV/√Hz.

The loop filter was designed such that the modulator remains stable up to over 50% modulation index (which according to DSD is then referred to as Full Scale). Due to the RTZ strategy, 100% modulation index corresponds to 3.75V peak or 2.65Vrms at the output of the D/A. Full scale inputs thus result in an output signal of 1.33Vrms. If there are no other contributing noise sources, this should result in an SNR of 120.5dB. We do these calculations, of course, to determine if this is indeed so after performing the measurements, and I shall keep you in suspense for a while.

The AD797 has an input-referred noise of 0.9nV√Hz, equal to 48Ω. Unfortunately, a 100Ω “stopper” resistor has to be inserted in the inverting input when operating the device at unity (noise) gain. Running through the calculation again yields 1463Ω or 5.1nV/√Hz, a noise contribution of -125dB.

9. PRACTICAL RESULTS

The converter was tied to the output of an AP2 test set and the data stream was recorded for FFT analysis. The clock source was a standard HCMOS type oscillator. A low-jitter part from Raltron had no discernable effect on the system performance. The pictures shown here are from the OPA627 version.
Signal-to-noise or unweighted dynamic range was measured by feeding the converter a -60dB input signal.
To extract “the number”, the energy in the FFT bins corresponding to the 1kHz input was integrated and similarly the energy in all other bins from 100Hz to 20kHz. Dividing the two and adding 60dB to relate to a full-scale signal yields the unweighted dynamic range equal to the signal-to-noise ratio. The resulting figures are:

SNR=DR=117dB (OPA627)
SNR=DR=120dB (AD797)

It may be surmised that another contributing noise source at -121dB is present, which is probably a combination of clock jitter and plain remaining quantisation noise.
Another important figure is full-scale THD+N. It is measured in the same way as DR, but then with a full scale signal. Any harmonic distortion and modulation noise will contribute to the measurement.
After trying to finger some unknown jitter source, the “skirts” around the 1kHz signal were found to be FFT window artefacts. The corresponding numbers are

THD+N=112dB (OPA627)
THD+N=115dB (AD797)

Why exactly the distortion was lower with the AD797 has not quite been determined yet.

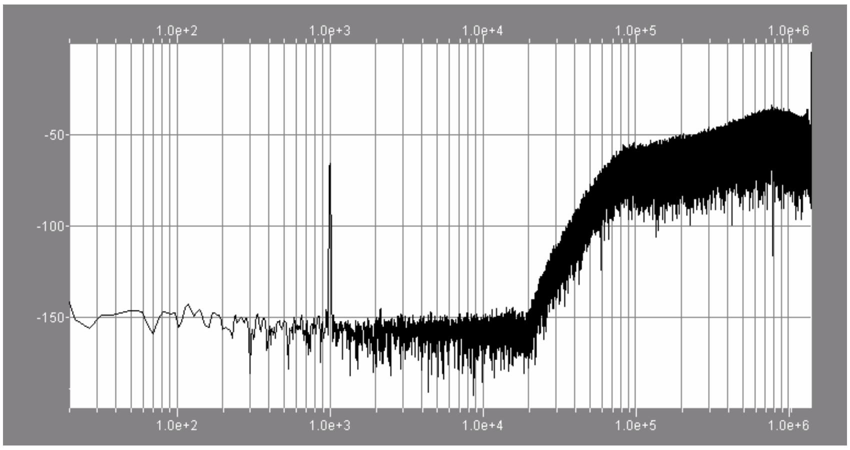

Measured output spectrum with a -60dB stimulus

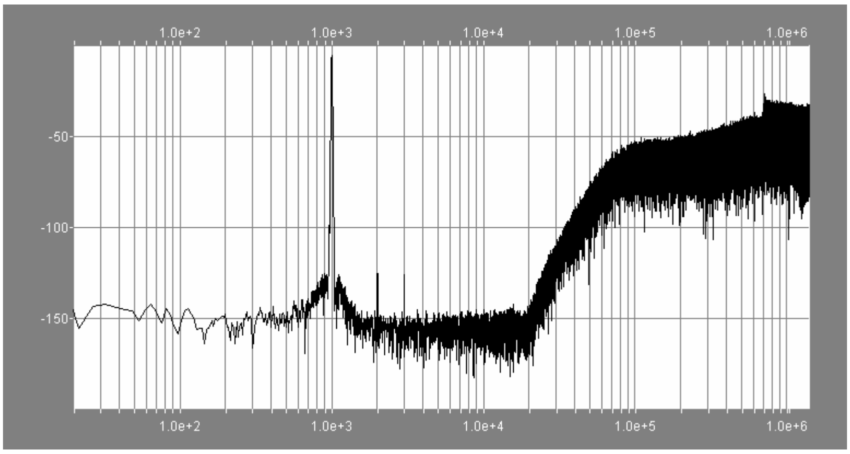

Measured output spectrum with a full-scale stimulus

10. FURTHER RESEARCH

Diligent work by others[2] has shown that one-bit noise shapers have some inherent evils lurking below the noise floor, which can be readily suppressed after converting to a multiquantum system. In this case this would entail using a PWM style modulator. While two-edged PWM is the easiest, one-sided PWM will probably be needed to prevent noise demodulation due to unequal switching transitions. This also opens up possibilities for increasing the sampling rate and hence the useful bandwidth, for increasing modulation index and for reducing remaining quantisation noise. It is thought practical to achieve a DR of 126dB and a flat noise floor up to 80kHz, which would make the circuit eminently suitable for 192kHz recording. The decision to downconvert to other release formats can then be made at the mastering stage.

11. CONCLUSION

Voltage switching A/D and D/A conversion was found, even with the simplest means to offer SNR performance on a par with the best charge switching converters and THD+N vastly superior to that of any other converter. The fact that such good converters can be built using standard discrete parts provides a solution to the dependence of the [high-end of the] pro-audio industry on semiconductor manufacturers' willingness to develop high-end parts for small production runs.

12. REFERENCES (samples)

[1] Steven R. Green et al., "An 18-bit Delta-Sigma D/A Processor System Achieving full-scale THD+N>100dB," presented at the AES 93rd convention, San Francisco, CA, 1992 October.

[2] Stanley Lipshitz and John Vanderkooy, "Towards a Better Understanding of 1-Bit Sigma-Delta Modulators," presented at the AES 112th convention, Amsterdam, The Netherlands, 2002 May 10-13.