An internally non-linear ADC for a $\Sigma\Delta$ accelerometer loop

Hanspeter Schmid[1], Sven Sigel[1], Marc Pastre[2], Maher Kayal[2], Pascal Zwahlen[3], Anne-Marie Nguyen[3]
[1]University of Applied Sciences NW Switzerland, Institute of Microelectronics; [2]Ecole Polytechnique Fédérale de Lausanne, Electronics Laboratory; [3]Colibrys SA, Neuchâtel, Switzerland. — Contact author: hanspeter.schmid@fhnw.ch

***Abstract*—If an analog-to-digital converter (ADC) is used within a $\Sigma\Delta$ converter such that the loop filter is mixed signal, then the ADC does not need to have the full resolution of the $\Sigma\Delta$ converter, and its input values will not be uniformly distributed over the input range. In this paper, we argue that the best input distribution to use as a general model is the maximum-entropy distribution (for location parameters: the Gaussian distribution). We then describe an externally linear ADC with non-linearly spaced decision boundaries implemented in 0.6 µm CMOS and show measurement results of a 19-bit dynamic range $\Sigma\Delta$ accelerometer loop inside which the ADC reaches more than the required 7-bit performance with only 2^6 decision levels.**

I. Introduction

One way to get a very good precision and linearity performance out of a sensor is to put it inside a $\Sigma\Delta$ loop [1], [2]. Fig. 1 shows a MEMS acceleration sensor inside such a loop. The sensor consists of a very small mass that is suspended between two plates by a spring. Any acceleration will change its speed and therefore its position and the capacitances to the top and bottom plate. The $\Sigma\Delta$ feedback in such systems is made by accelerating the mass with an electrostatic feedback force [1]. In order to make the system adaptable to different types of sensors, the sensor read-out signal is converted into the digital domain immediately after the sensor read-out amplifier such that the $\Sigma\Delta$ loop filter can be digital and programmable.

Since there are already two mechanical integrations happening before the signal comes into the electrical domain, the digital filter needs only be of third order to make a fifth-order $\Sigma\Delta$ loop. The first integrator in a $\Sigma\Delta$ loop has the strongest specs on noise and linearity, and the specs of the following integrators are relaxed [3], therefore the 19-bit dynamic range necessary for inertial navigation applications can be obtained with only 7 bits of resolution in the ADC [1].

This does, however, not mean that it is necessary to build a linear 7-bit ADC. The probability distribution at the input of the ADC is not uniform, in most cases there are more values to processed in the centre of the ADC range than towards the extremes. This fact has been used before in multi-bit $\Sigma\Delta$ converters by using complex µ-law ADCs [4] or — easier to implement — semi-uniform quantisers [5].

Our method differs from earlier publications because we outline a theoretical basis for assuming a probability distribution for such a system. First, we argue why we should assume a Gaussian distribution of ADC input values. Then we derive the quantisation intervals and the quanta for the optimum quantiser, present a performance simulation, discuss a 0.6 µm CMOS implementation, and then present measurement results.

II. ADC input probability distribution

Neither [4] nor [5] explained why the specific nonlinearities were chosen. According to [6], a quantiser can be made optimum if the probability density of its input values is known. The question is now: which probability distribution shall we choose if we do not know much?

Simulations and measurements of an acceleration sensor loop have shown that the specific distribution depends on the signal at the input, the sensor, and the parasitic effects like offset and electronic noise. Four histograms of simulated and measured values are shown in Fig. 2; not all of them look Gaussian.

A $\Sigma\Delta$ accelerometer loop — like most other $\Sigma\Delta$ sensor loops too — uses high-gain feedback to keep the sensor mass at a given centre position. The sensor displacement is then small compared to the maximum possible displacement before the sensor mass touches the top or bottom plate. However, since the distributions look so different for different input signals, we have no further knowledge about the probability distribution except that we measure a location parameter with hard bounds far away from the operating range.

Therefore, the distribution to assume is the one which makes *no implicit assumptions other than that we measure a location parameter*. This is the so-called maximum-entropy distribution for location parameters, which is the Gaussian distribution [7].

So we do not choose the Gaussian distribution because we know that the values to be measured will have that distribution — they often have not — but because it it is the best representation of our prior knowledge: the user of our system can attach different sensors, do different applications, have different input signals, and we do not know what probability distributions to expect, just that we measure is a location parameter. In the remainder of this paper, we focus on an acceleration sensor, but this reasoning extends to other $\Sigma\Delta$ systems as well because the output of the second (or even the first) integrator also are of location parameter nature.

III. 64-level externally linear ADC

In general, a quantiser has n quantisation intervals Q_i with boundaries x_i

$$Q_i = \{x : x_{i-1} < x \leq x_i\} \quad \text{for } 0 < i < n \tag{1}$$

where an input value x lying inside the quantisation interval Q_i is represented by an output value $y = q_i$ such that

$$x_0 < q_1 < x_1 < q_2 < x_2 < \ldots < x_{n-1} < q_n < x_n \tag{2}$$

978-1-4244-5309-2/10/$26.00 ©2010 IEEE

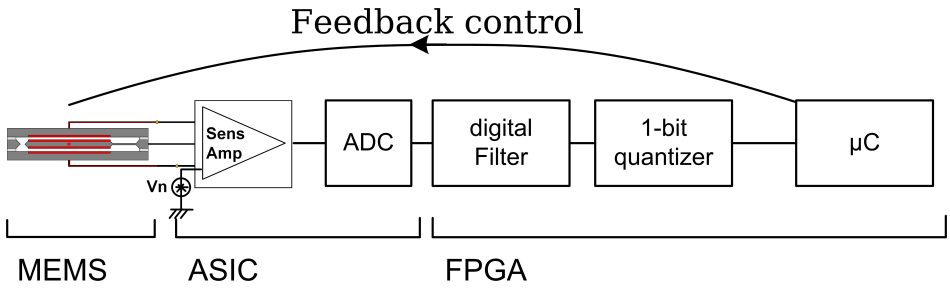

Fig. 1. Block diagram of a $\Sigma\Delta$ accelerometer loop.

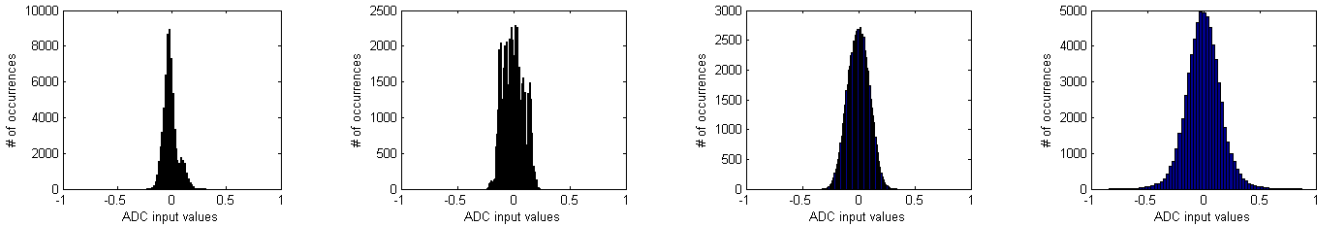

Fig. 2. Histograms of the signals at the ADC input. From the left: measured with 1g DC | simulated 0g 100Hz | 4g 100HZ | 11g 100Hz.

and $x_{0,n} \to \pm\infty$. The mean square error of the quantisation noise (the quantisation noise power) then becomes

$$P_N\left(q_i, x_i\right) = \sum_{i=1}^{n} \int_{x_{i-1}}^{x_i} \left(q_i - x\right)^2 f(x) dx \tag{3}$$

where $f(x)$ is the probability density of the inputs of the quantiser. P_N needs to be minimised.

It was shown in [6] that the best quanta q_i for given boundaries x_i are the centres of mass of the probability density function $f(x)$ in the respective range:

$$q_i = \frac{\int_{x_{i-1}}^{x_i} x \cdot f(x) dx}{\int_{x_{i-1}}^{x_i} f(x) dx} \,. \tag{4}$$

On the other hand, the best partition points x_i for given q_i are

$$x_i = \frac{1}{2}\left(q_i + q_{i+1}\right) \quad \text{for } i = 1 \ldots n-1 \,. \tag{5}$$

It is possible to use any starting point and use (4) and (5) iteratively to derive the optimum x_i and q_i. Such an iteration would, however, only find a local minimum, so it is important to have a good starting point.

To find that point, we used an approximating method from [6]: the ideal quantiser would have Q_i such that the area under $f(x)$ is the same for all Q_i, and since the integral over a probability density $f(x)$ is always 1, this area is $1/n$. Therefore, the x_i are the solutions of the integral equations

$$\int_{-\infty}^{x_i} g(x) dx = \frac{i}{n} \,, \tag{6}$$

where $g(x)$ is the asymptotic density of quanta, i.e., the density of quanta of an optimum quantiser with $n \to \infty$. Using Taylor approximations, it can be shown that for a given $f(x)$, this asymptotic density can be obtained as

$$g(x) \approx \frac{f^{\frac{1}{3}}(x)}{\int_{-\infty}^{\infty} f^{\frac{1}{3}}(x) dx} \,; \tag{7}$$

while this approximation does not give the optimum solution, it gives a solution quite close to the optimum.

So we defined $f(x)$ as a Gaussian distribution with mean 0 and standard deviation 1 and calculated $g(x)$ according to (7). Then we calculated the boundaries of the Q_i with (6). This we used as a starting point for the iteration described by (5) and (4). Finally, we scaled the results to adapt the x_i to the voltage range of the ADC in our system, $x_{1,n-1} = \pm 0.9\,\mathrm{V}$, and the q_i to a normalised digital range, $q_{1,n} = \pm 1$.

The smallest quantisation interval of our ADC has the width $x_{32} - x_{31} = 15.4\,\mathrm{mV}$. This is even larger than the interval of a linear 7-bit quantiser, $1.8\,\mathrm{V}/2^7 = 14.1\,\mathrm{mV}$, as it would be needed to build a semi-uniform quantiser like the one in [5], and it is much larger than the very small quantisation steps used in the centre range of µ-law ADCs [4].

Note that this quantiser is built for an input distribution that has a specific standard deviation. The stage preceding it must therefore be a programmable-gain amplifier whose gain can be set differently for different applications and sensors. In our system, it is programmable in the range $10 \ldots 230$.

The schematic of our ADC is shown in Fig. 3. It consist of a block containing comparators and a resistive divider, similar to a flash ADC, but constructed such that the decision thresholds are the x_i. The output is a thermometric code from which the number i of the quantisation interval is calculated. Finally, a look-up table (in our test setup residing on an FPGA) assigns the q_i for each Q_i with a precision of 16 bits.

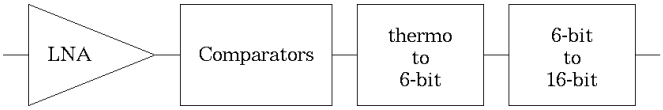

Fig. 3. Schematic of the externally linear internally non-linear ADC.

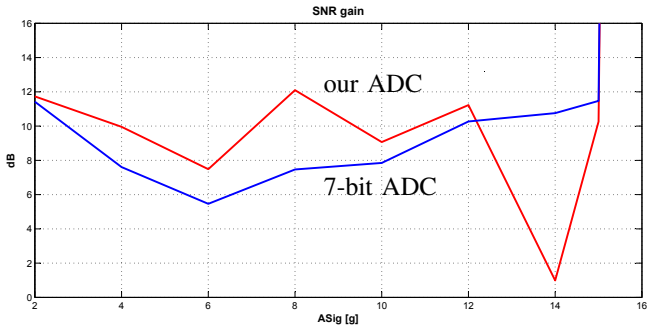

Fig. 4. Comparison of the increase in SNR that a linear 7-bit ADC gives compared to a linear 6-bit ADC (bottom curve) and that our ADC gives compared to a linear 6-bit ADC (top curve).

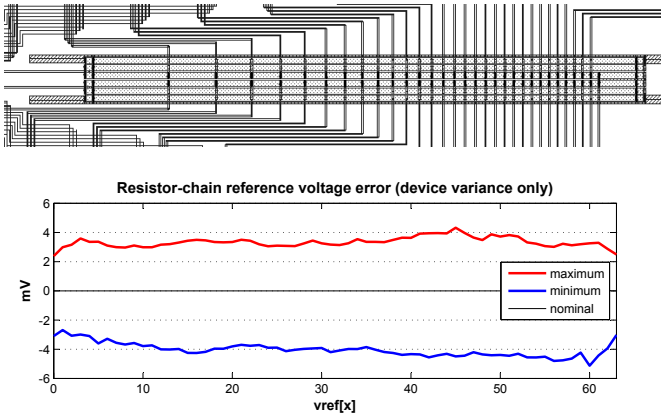

Fig. 5. Layout of the resistor chain with horizontal resistor elements and vertical metal lines to access them, including dummies (top) and worst-case errors on the different x_i obtained with 1000 Monte-Carlo runs (bottom).

IV. Simulation

We simulated such a system with an acceleration sensor; the system model comprises the non-linearities of the sensor as well as the electronic noise of the switched-capacitor LNA and of the reference voltages.

Our simulation model neglects all mechanical resonances of the sensor except the first one, so we cannot be sure that we simulate the absolute value of the shaped quantisation noise correctly. However, the ratio between simulations with the same sensor but different ADCs will be accurate. Therefore we simulated the same system with a linear 6-bit ADC, a linear 7-bit ADC, and our internally nonlinear ADC.

Figure 4 shows the increase of SNR when going from a 6-bit to a 7-bit ADC and when going from a 6-bit to our ADC. It shows that our ADC performs better than the linear 7-bit ADC for all signal magnitudes except for the highest one. So we can conclude that by allowing a slightly lower SNR for the maximum input signal (which is definitely no problem in an accelerometer), we can achieve a better performance with our ADC that would be achieved with a linear 7-bit ADC.

V. Design

The ADC has been implemented on a two-metal 0.6 µm process along the lines of a conventional flash ADC [8]. The resistive divider was built with resistors of different sizes such that the x_i could be implemented as precisely as possible; the different lengths of the resistors in the one-time-folded resistor chain with dummies give the beautiful tapered structure in the layout shown in Fig. 5 (top).

Since a small scaling error does not matter within a $\Sigma\Delta$ converter — the output magnitude is set solely by the feedback, and not by the magnitude in front of the 1-bit quantiser — we decided to connect the resistive divider between the (regulated) 3.3 V supply of the chip and ground, without a dedicated regulator setting a precise reference. The simulated errors on the x_i are shown in Fig. 5 (bottom).

Two types of comparators were designed, standard two-stage comparators [8], with pMOS diff-pairs for the x_i below zero and nMOS diff-pairs for the other x_i. Their random offset voltage was made small enough compared to the smallest quantisation step. So our ADC has the same number of comparators as a linear 6-bit ADC, but the comparators have the size and power consumption as if they were used in a linear 7-bit ADC. So our ADC, while giving even a bit better performance *in this* $\Sigma\Delta$ *sensor application* than a linear 7-bit ADC, is half as large, has half the input capacitance, and consumes half as much power than a linear 7-bit ADC would.

Finally, the thermometric-to-6-bit encoder could be realised either as a bubble remover encoder, or as a top-hot encoder, or simply as a 64-input adder. On the 0.6 µm process we used, the adder would be fastest with a delay of 30 ns, but also the largest with 210'000 µm^2 active area. The top-hot encoder only has 80'000 µm^2 but requires 55 ns to encode. The bubble remover also takes 55 ns and requires 120'000 µm^2.

The constraints of the full system were that a voltage on the sensor can be sensed, processed by a correlated-double-sampling amplifier, and then converted into the digital domain, such that the last two steps are finished within 500 ns. To save power, we decided to construct a free-running ADC; e.g., the last amplifier stage is connected directly to the flash converter stage that is also connected directly to the encoder (with no sampling in between), such that all three settle together, and in the end the digital output of the encoder is sampled. The worst-case total settling time over the whole chain then needs to be 500 ns. Having the speed of a full adder was not necessary, and since this kind of flash ADC should not give bubbles in the thermometric code, we chose the top-hot encoder for the system in order to save chip area. The simulated performance of the ADC by itself is given in Table I.

Power consumption	1.33	mW
Conversion time (incl. pre-amp/encoder)	500	ns
Resolution (in application)	> 7	bits
Layout area flash ADC	440'000	μm^2
Layout area encoder	310'000	μm^2
Input capacitance	15	pF

TABLE I
SIMULATED WORST-CASE PERFORMANCE OVER PROCESS CORNERS AND A $-40\ldots+85°$C TEMPERATURE RANGE (WORST CASES FROM CORNER AND MONTE-CARLO SIMULATIONS).

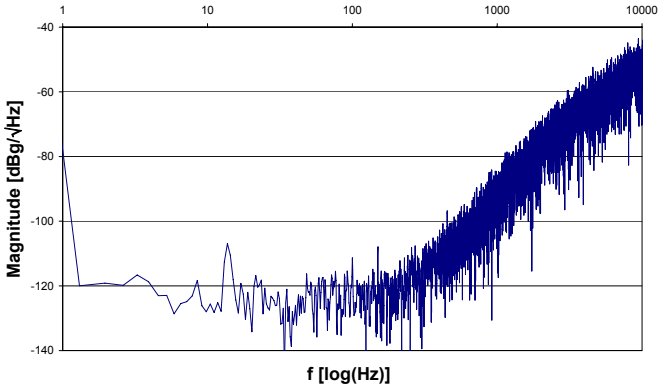

Fig. 6. Example for a measured power spectrum.

VI. MEASUREMENT RESULTS INSIDE THE SYSTEM

We will now not present a conventional measurement of this ADC, as the point of this is that it performs properly *within a sensor system*. The main point of this paper is the idea is that it is the Gaussian distribution that should be used, and this can only be validated by making lots of measurements with different sensors under different conditions and with different input signals. These have been done: our test board allowed to switch back and forth between our ADC (on chip) and a commercial 12-bit linear ADC (off chip) of which only the six or seven MSBs were used to simulate 6-bit and 7-bit ADCs; we made extensive measurements with real sensors and sinusoidal inputs on a shaker table and found no system state in which our ADC was responsible for poor system performance. The full-system SNR performance for one type of sensor has been discussed in detail in [1]; measurements of linearity and long-term stability were shown in [9].

Here we can only give one example of the type of measurements performed to confirm that our ADC works inside the applications. Fig. 6 shows a power spectrum measured with a sensor attached to the test setup, which was located in the basement of a building where gravitational noise (i.e., building vibrations) were minimum. We chose the sensor, filters and actuation voltages such that the noise of the ASIC would dominate. The measurement gave a noise density of $1.06\,\mu g/\sqrt{\text{Hz}}$. Calculations with a 7-bit ADC basing on a linearised model of the $\Sigma\Delta$ loop give $1.36\,\mu g/\sqrt{\text{Hz}}$, which is 0.36 bit higher, indicating an ADC performance *in this particular measurement* of 7.36 bit, which agrees with Fig. 4.

VII. CONCLUSION

We have shown how to replace the linear ADC in a mixed analog-digital $\Sigma\Delta$ sensor loop by an externally linear, internally nonlinear ADC which reaches a bit more than 7-bit performance with only 2^6 comparators. This result extends to any $\Sigma\Delta$ system in which an ADC is needed because a part of the loop is digital, and it also extends to constructing multi-bit quantisers for $\Sigma\Delta$ converters as has been done in [4], [5].

Although the design was made on a 3.3-V two-metal 0.6 µm process, targeting low-volume production, the method scales well to all processes where there still is enough supply voltage to build a conventional flash ADC.

The design of the ADC is straightforward, but the validation is not: the quantisation noise of an ADC depends on the probability distribution at the input, and this depends on sensor physics filter parameters, and the applied signal. So unless we test *all* possible sensors and signal classes, we will not know for sure precisely how the performance of our ADC compares to the performance of a linear ADC. However: the quantisation noise of a linearly spaced ADC also depends on the probability distribution at its input [3], so it is just as difficult to validate the performance of a conventional ADC in a closed $\Sigma\Delta$ sensor loop. As we have not seen a case in our extensive simulations and measurements in which our non-linear ADC failed, we conclude that it is as safe to use as a linear ADC.

This result has impressed us; the theory we used to derive a probability distribution at the ADC input only guarantees that the ADC we build best reflects the information we have about such systems, but not that the information we have is sufficient. So it was a pleasant surprise that our simulations and measurements confirmed that our approach is useful.

VIII. ACKNOWLEDGEMENTS

We would like to thank the Swiss Innovation Promotion Agency CTI for the financial support (project 7472.2 NMPP-NM), the Cadence Academic Network for their continuing support with tools and methods, and Alex Huber for many interesting discussions about the contents of this paper.

REFERENCES

[1] M. Pastre, M. Kayal, H. Schmid, A. Huber, P. Zwahlen, A.-M. Nguyen, and Y. Dong, "A 300Hz 19b DR capacitive accelerometer based on a versatile front end in a 5th-order delta sigma loop," in *Proc. ESSCIRC*, Athens, Greece, Sep. 2009, pp. 288–291.

[2] C. Condemine, N. Delorme, J. Soen, J. Durupt, J.-P. Blanc, M. Belleville, and A.Besançon-Voda, "A 0.8 mA 50 Hz 15b SNDR $\Delta\Sigma$ closed-loop 10g accelerometer using an 8th-order digital compensator," in *Proc. ISSCC*, San Francisco, 2005, pp. 248–249.

[3] R. Schreier and G. C. Temes, *Understanding Delta-Sigma Data Converters*. IEEE Press, 2005.

[4] Z. Zhang and G. Temes, "Multibit oversampled $\Sigma\Delta$ A/D convertor with nonuniform quantisation," *Electronics Letters*, vol. 27, no. 6, pp. 528–529, Mar. 1991.

[5] B. Li and H. Tenhunen, "Sigma delta modulators using semi-uniform quantizers," in *Proc. ISCAS*, vol. 1, Sydney, May 2001, pp. 456–459.

[6] S. P. Lloyd, "Least squares quantization in PCM," *IEEE Trans. Inf. Theory*, vol. 28, no. 2, pp. 129–137, Mar. 1982.

[7] E. T. Jaynes, *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.

[8] D. A. Johns and K. Martin, *Analog Integrated Circuit Design*. New York: John Wiley & Sons, 1997.

[9] P. Zwahlen, A.-M. Nguyen, Y. Dong, F. Rudolf, M. Pastre, and H. Schmid, "Navigation grade MEMS accelerometer," in *Proc. IEEE MEMS Conference*, Hong Kong, Jan. 2010.