Published in International Symposium on Low Power Electronics and Design, 247-252, 1996
which should be used for any reference to this work

Implementation of a Micro Power 15-bit 'Floating-Point' A/D Converter

L. Grisoni, A. Heubi, P. Balsiger and F. Pellandini
Institute of Microtechnology University of Neuchâtel
Tivoli 28, CH-2003 Neuchâtel-Serrières, Switzerland
Phone: # 41 38 30 16 31, Fax: # 41 38 30 18 45, E-Mail: grisoni@imt.unine.ch

ABSTRACT

Micro power A/D converter are required for power sensitive, battery-operated equipment such as hearing aids. This paper overviews the principles of the 15-bit 'Floating point' converter and presents its implementation in a low voltage 2 µm CMOS technology. The die area is 1.4 by 1.4 mm and the power consumption 50 µW at ±1.25 V and 16 kHz sampling frequency. Measurement showed that the internal noise level is higher then expected resulting in a reduced dynamic range of 13 to 14 bits. Informal listening tests showed a very good speech quality.

1. INTRODUCTION

Progress in low-power microelectronics technology and digital signal processing has opened the way to numerous digital portable applications. In particular, new audio applications have been developed and exhibit a very fast growing market. Low power A/D converters targeted for audio signals are thus required to reduce the overall power consumption and consequently the battery life.

Typically, for hearing aid devices, the audio signals cover a dynamic range of about 80 dB and to ensure satisfactory quality a bandwidth of 8 kHz must be provided. Though, because of masking properties of the human ear, the requirement on the noise level (SNR) is only about 30 dB [4].

The idea of a ‘floating point’ converter which separates dynamic range and resolution to decrease power consumption was presented in 1992 [1]. This first design, however, did not fully meet the expectation when dealing with demanding audio signals. The concept was thus improved [3] resulting in a new 15-bit, 9 bits mantissa floating point A/D converter.

This paper focuses on the design and implementation of the above converter in a low voltage CMOS technology. Section 2 briefly reviews the revised ‘floating point’ A/D converter while section 3, 4 and 5 deal with hardware realization. Section 6 presents the implementation results and section 7 concludes the discussion.

2. FLOATING POINT A/D CONVERTER PRINCIPLES

A concept for a “floating point” or “relative precision” A/D converter (figure 2.1) was first published by A. Schaub [1] in 1992. The main idea was to adaptively scale the input signal in such a way that most of the time, it would fit well to the fixed conversion range of a rather coarse quantizer. Scaling gain and quantizer output were subsequently combined in order to properly represent the signal sample within the full dynamic range of the acoustic input signal.

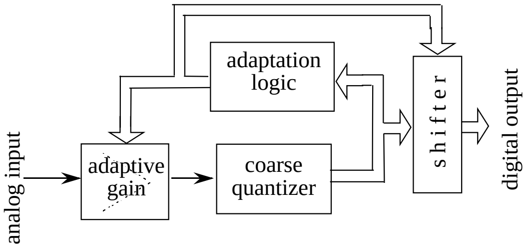

Fig. 2.1: *Block diagram of floating point A/D*

More precisely, in Schaub’s converter, the signal is fed to a programmable amplifier with gain settings ranging from 0 to 46.5 dB in steps of 1.5 dB. They are set by an adaptation logic and the amplified signal is applied to a 6-bit quantizer (5-bit magnitude plus sign).

The five bits representing the magnitude of the quantized sample serve as the input to the adaptation logic which updates the gain of the programmable amplifier according to an adaptation table which is inspired by N.S. Jayant’s paper [2] on adaptive quantization of PCM signals. The multipliers are restricted to powers of $2^{1/4}$, corresponding to approximately 1.5 dB steps and the target value is half the maximum amplitude range.

The performance of this converter was evaluated through simulations using male and female speech as well as music inputs. The input level of the test signals was adjusted to -30 dB with respect to the full range of the converter, corresponding to realistic conditions in audio applications such as hearing aids. Typical SNR values around 26 dB resulted from these experiments.

Comparative listening tests using headphones still revealed noticeable differences between original and processed signals. Even though the amplified signal exceeded the fixed quantizer range in only a small fraction of time, the resulting clipping effect was clearly perceived.

Improving the signal to quantization noise ratio seemed straightforward at first glance. In a series of experiments the number of bits in the quantizer was repeatedly increased and the adaptation table was extended as required. The performance of the converter was observed by means of segmental SNR values computed for subsequent time intervals of 20 ms each. As long as the amplified input signal was within the range of the quantizer, improved SNR values were observed as expected. When the maximum amplitude of the quantizer was exceeded, however, the additional quantizer bits failed to improve the segmental SNR.

These effects are illustrated in figure 2.2. At the top a speech segment of 1 s duration is shown. Segmental SNR curves for Schaub's converter with an increasing number of quantizer bits are shown in the middle.

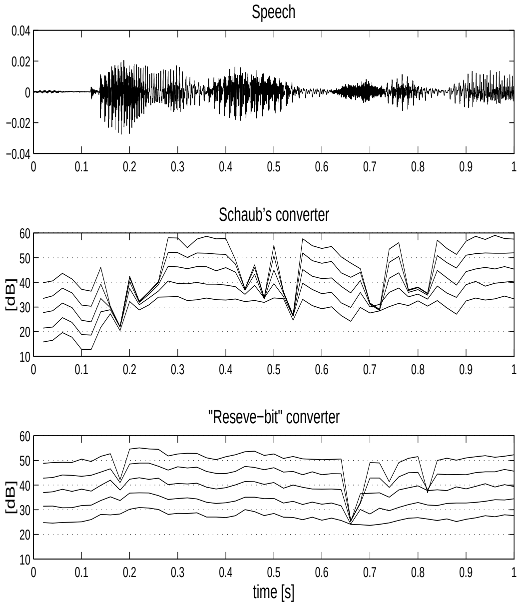

Fig. 2.2: *Segmental SNR for 6 to 10-bit quantizers*

A new series of experiments was then started with the concept of a so-called "reserve bit" [3]. For a quantizer with a given number of bits, the idea was to reduce the target value from half to a quarter of the maximum amplitude. In this way, a larger margin to protect against clipping was introduced. Segmental SNR values as well as comparative listening tests confirmed the expected improvement which are depicted at the bottom of figure 2.2.

Since the intention is to develop a self-contained converter unit a time-varying alignment of the quantizer codeword to the output word in accordance with the applied gain in the input amplifier is required. Nevertheless, a on-chip multiplier is out of scope since it would result in a much higher complexity and power consumption.

In his former design, Schaub proposed to use look-up tables with pre-multiplied codewords. Unfortunately, the size of the look-up tables grows exponentially with the number of quantizer bits. Therefore, this approach is not very attractive in the case of quantizers with a higher number of bits.

In the final design of the revised "relative precision" converter [3] the input gain is computed in 1.5 dB steps though the input amplifier changes are brought to effect only when they have accumulated to 6, -6, or -12 dB. Hence, shifting operations are sufficient for codeword alignment at the back end.

Simulations showed that segmental SNR values dropped by only 1 dB typically compared to the case where the amplifier gain is actually changed in 1.5 dB steps. Obviously, the previously introduced "reserve bit" was quite helpful in this situation as well. To be precise, the performance curves at the bottom of figure 2.2 actually reflect the converter design with the coarse 6 dB gain adjustments already taken into account.

At the end of a series of comparative listening tests it was decided to implement a floating point converter with a 9-bit quantizer and a maximum gain of 36 dB, thus yielding a dynamic range of a 15-bit.

3. ADAPTIVE GAIN

The 'reserve bit' adaptation table for a 9 bit quantizer is presented on table 3.1 which shows both the increment / decrement values in dB and its two's complement representation. The magnitudes of column one are obtained from the six most significant bits of the coarse quantizer. The gain computation can thus take place as soon as the 6th bit is known. This time constraint as well as the seven gain values (0, 6, ..., 36 dB) are the starting point to the design of the adaptive gain.

The adaptive gain can be realized very simply with an inverting amplifier as described in figure 3.1. The gain value is given by the ratio of the total feed-back capacitor and C1. Such a structure uses a single active element and benefits from the good capacitor matching of CMOS technologies. Switches S1 to S7 are used to obtain the desired gain value while Sr1 and Sr2 guarantee a proper reset of the capacitors. A resistor R is used to reduce the bandwidth and consequently the noise level on C1.

coarse quantizer		increment/decrement	
magni.	$b_5b_4b_3b_2b_1b_0$	dB	$\partial_3\partial_2\partial_1\partial_0$
0 .. 63	0 0 0 X X X 1 1 1 X X X	+ 1.5	0 0 0 1
64 .. 71	0 0 1 0 0 0 1 1 0 1 1 1	- 1.5	1 1 1 1
72 .. 79	0 0 1 0 0 1 1 1 0 1 1 0	- 3	1 1 1 0
80 .. 95	0 0 1 0 1 X 1 1 0 1 0 X	- 4.5	1 1 0 1
96 .. 111	0 0 1 1 0 X 1 1 0 0 1 X	- 6	1 1 0 0
112..119	0 0 1 1 1 0 1 1 0 0 0 1	- 7.5	1 0 1 1
120..255	0 0 1 1 1 1 1 1 0 0 0 0 0 1 X X X X 1 0 X X X X	- 9	1 0 1 0

Table 3.1: *LC1 truth table*

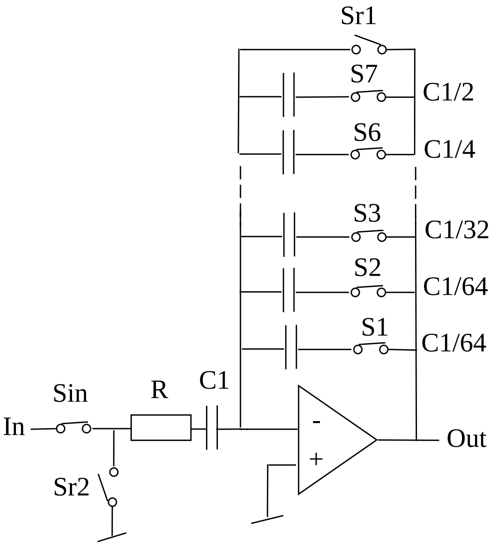

Figure 3.1: *Schematic of adaptive gain*

The scheduling of the adaptive gain is presented on figure 3.2. 'Bit #' indicates the bit presently computed in the coarse quantizer and its basic clock 'Ck'. 'A' is the moment at which the sixth bit is known and thus adaptation logic can start computing the control signals of S1 to S7 which are then set at 'B'. 'C' is the beginning of the core cycle i.e. when the analog signal to be converted is loaded on the input capacitor of the coarse quantizer. Two phases can thus be acknowledged: i) Sin is open while all the other switches are closed to ensure proper reset; ii) amplification takes place and Sin is closed while S1 to S7 are set according to the adaptation logic.

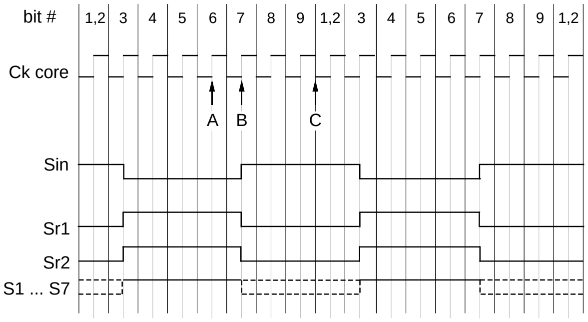

Figure 3.2: *Adaptive gain scheduling.*

C1 is chosen in such a way that the sampled noise remains smaller then half a LSB. The maximal value of R is defined by the load characteristic of C1 and ensures a load error smaller then half a LSB at switching time C (see figure 3.2).

The active element is a two stages OTA as depicted in figure 3.3. It ensures both high output swing and relatively high DC gain while keeping the power consumption low. M7, M8 and C_c are necessary to improve the stability.

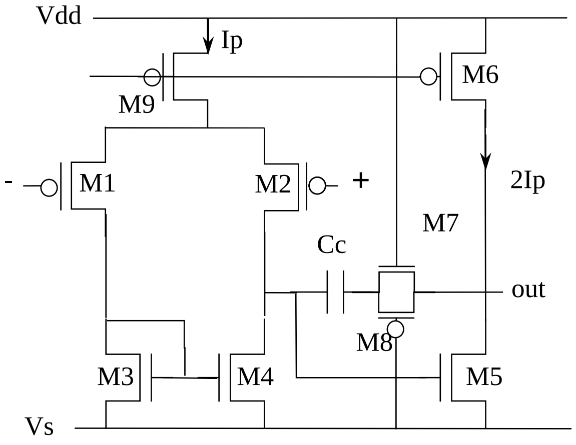

Figure 3.3: *Schematic of OTA*

The specifications to design the OTA are the DC gain (> 60 dB) and the equivalent input (< 50 µV). This value depends on the input stage, which in the considered application consisted in a microphone and a pre-amplifier with a total noise level of 100 µV. Of course minimum power consumption is desired.

Using a 1.4 µA current results in an OTA which characteristics are given in table 3.2. The achieved power consumption is less then 16 µW at ±1.3V and the phase margin is just sufficient.

A simple R-based, weak inversion current source [7] is used to polarize the OTA. The obtained polarization current is proportional to U_T and consequently the transconductance G_m of the differential pair of the OTA is independent of the temperature. Furthermore, if a poly resistor is used, the current becomes independent, in first approximation, of the supply voltage as well ($\delta R/\delta Vdd=0$). The draw back though is the poor resistor precision resulting in poor Ip value precision.

open loop gain	66.8 dB
gain bandwidth	4.6 e+06 Hz
phase margin	26°
offset	0,0003 V
output swing low	-1.247 V
output swing high	1.29 V
input noise	4.6 e-08 V/Hz$^{1/2}$
slew-rate	3 e+06 V/s
power dissipation (at ± 1.3V)	15.4 e-06 W

Table 3.2*: Performances of designed OTA*

The effective behavior of the adaptive gain is determined by some parasitic effects such as finite DC gain and offset of the OTA, leakage current and charge injection due to the switches etc... Those effects are not easily numbered but a qualitative study is possible.

The finite DC gain results in smaller amplification values. However, the shifter following the coarse quantizer performs exact division by 2, 4, etc. Audio simulations of the whole floating point A/D, with real gain values obtained from ELDO simulations and exact divisions showed that the speech quality was considerably decreased. Experiments lead to the conclusion that gain errors of less then 1% were required to maintain a good quality. A simple solution to compensate for the finite DC gain is to slightly modify the feed back capacitor values in such a way that a tolerable gain value is reached.

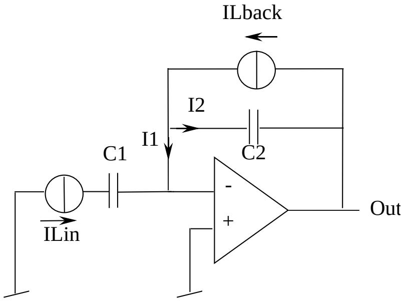

Figure 3.4: *Model for leakage currents*

Leakage currents due to switches must be considered and the model of figure 3.4 can be used. I_{Lin} and I_{Lback} are the sum of the leakage currents of Sin and Sr2 and all feed back switches respectively. The resulting tension on the capacitors are given by equation 3.1. $dU_{2,1}$ is the tension that should be applied on C1 which would result in dU_2 on C2. The global effect of the leakage currents, analyzed on C1 is thus given by 3.3.

$$dU_1 = \frac{I_1 \cdot dt}{C_1} \qquad dU_2 = \frac{I_2 \cdot dt}{C_2} \qquad dU_3 = \frac{I_{Lin} \cdot dt}{C_1} \qquad (3.1)$$

$$dU_{2,1} = \frac{I_2 \cdot dt}{C_2} \cdot \frac{C_2}{C_1} \qquad (3.2)$$

$$dU_{in} = dU_1 + dU_{2,1} + dU_3 = \frac{I_{Lin} + I_{Lback}}{C_1} \cdot dt \qquad (3.3)$$

I_{Lback} and I_{Lin} can legitimately be assumed constant as well as dt which is the load time shown in figure 4.2 (B to C). The overall result is thus an input offset, which is independent of the feed back capacitor value C_2.

Another important parasitic effect is the charge injection due to the opening of Sr_1 and Sr_2. Each time Sr_1 opens, injection takes place resulting in an voltage offset. The latter is nevertheless constant since the total capacitance on the negative input of the OTA does not depend on the on/off state of S1 to S7. When Sr_2 opens, charge injection occurs on C_1.

Parasitic capacitors have been neglected in all the above considerations. Their effects are complex to evaluate but one can assume that they contribute to the offsets and modify the actual gain values. The offset can easily be taken care of in the DSP architecture that follows the A/D converter while gain values can be corrected by modifying the feed back capacitor.

4. COARSE QUANTIZER

Alexandre Heubi, from IMT, proposed in [5] a very simple switched capacitor circuit which performs RSD [6] cyclic conversion.

The accuracy of this type of converter is limited by a few factors such as offset due to the active element and charge injection. The finite gain as well as capacitor non-linearity and intrinsic noise also influence the accuracy. Finally the capacitor mismatch limits the maximum achievable SNR. Though this structure has very interesting advantages: the constant offset can be digitally corrected easily, the single active element is well suited for low power considerations and since capacitor matching is of some ‰, a SNR of about 60 dB can be reached.

The floating point A/D requires a coarse quantizer of 9 bits which is realized as a 8 cycles RSD converter. The converter presented in [5] can thus be redesigned with lightened specifications.

Sampled thermal noise, load characteristic and minimal area result in capacitors and resistors of 1.3 pF and 100 kΩ respectively.

Spectral input noise density of 9e-7 V/Hz$^{0.5}$ and minimal slew rate (SR) of 1e6 V/s define the transistors sizes and polarization current of the OTA. This ensure that the load is not critically limited by SR and that the accumulated input noise level, when computing the last bit, is sufficiently low for the comparators to take the right decision. A polarization current of 1 µA is used, leading to a OTA power consumption of 16 µW (including control logic) at ± 1.3 V.

5. ADAPTATION LOGIC

The block diagram of the adaptation logic is presented in figure 5.1. The six most significant bits of the coarse quantizer are fed to a combinational logic LC1 which compute four bits to code the 1.5 dB increment or decrement. The latter or an external value Gin depending on CTR, is used by the next block to calculate and memorize the new gain. Three bits are used to code the seven (1, 2, 4, ...) possible values which are then fed to a second combinational logic LC2 where seven control signals are generated. for the N respectively P MOS switches of the adaptive gain stage The gain is also fed to the shifter for memorization and proper shifting of the coarse quantizer's output.

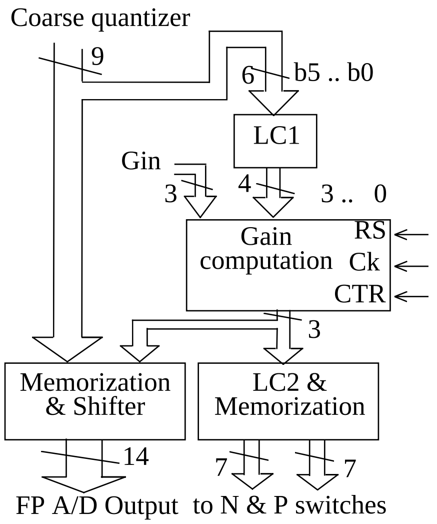

Figure 5.1: *Adaptation logic*

LC1 is described by table 3.1 and figure 5.2 gives the schematic of the gain computation . The upper part accumulates and stores the gain while the central one takes care of boundary conditions (minimum and maximum gain value). The lower part provides the necessary hardware to externally control the gain value (used for debugging purpose only). LC2 is derived from figure 3.1 to ensure proper switching for the desired gain.

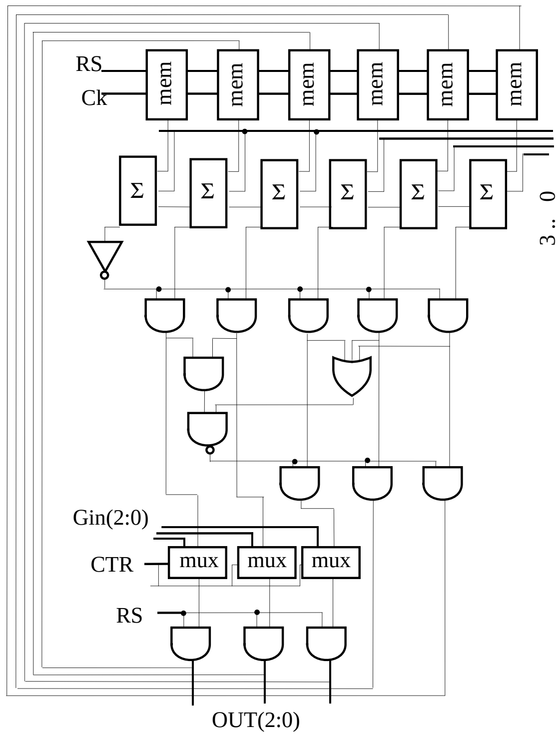

Figure 5.2: *schematic for gain computation*

6. IMPLEMENTATION

The 15-bit floating point A/D converter was implemented in the ALP2 LV double metal double poly 2 µm technology from EM Microélectronique Marin. Standard cells CSEL_LIB from CSEM and full custom layout for the analog parts were used. Extra multiplexers and control pins were added to allow full characterization and study of the obtained chip. The total area is about 2.5 by 2.5 mm with a die size of 1.4 by 1.4 mm. The layout is given in figure 6.1 where the upper section is the adaptation and control logic (≈2900 transistors) while the lower section contains the adaptive gain (left) and coarse quantizer (middle).

The measured power consumption is 50 µW at ±1.25V and 16 kHz. Informal listening test showed very good audio quality for speech and music. Nevertheless, silences are tarnished by noise. This is a common problem in AGC where the adaptive gain ceils at the highest gain resulting in maximum noise when the input signal is close to null.

Measurement of the offsets showed that the gain values are not perfectly doubled but remain in the error tolerance. As expected the behavior of the adaptive gain is very sensitive to the polarization of its OTA. The noise level is also slightly higher then foreseen resulting in a reduced dynamic range of about 80 dB.

The coarse quantizer works perfectly and meets its specifications.

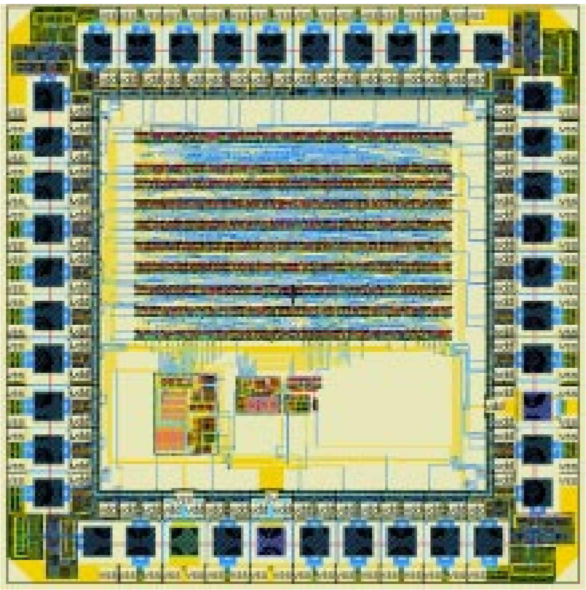

Fig: 6.1: *Layout of 15-bit floating point A/D converter*

Figure 6.2 shows the frequency response of the converter to a 500 Hz -9 dB sinus input signal. The y axis shows the level below full scale in dB while the x axis shows the frequency in Hz. A SNR of about 50 dB is reached. This can be explained because of the adaptation logic which targets one fourth of the coarse quantizer i.e. 7 bits.

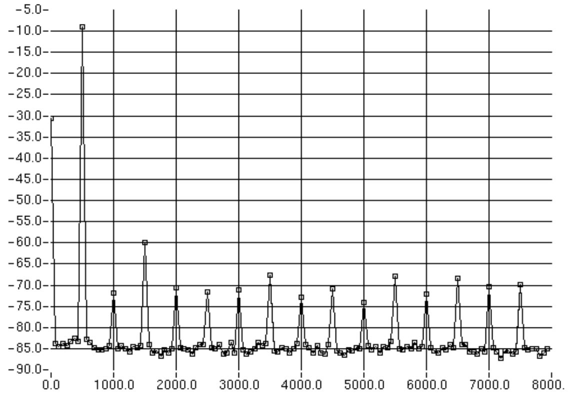

Figure 6.2: *Frequency response of the converter*

7. CONCLUSIONS

To decrease power consumption, the floating point A/D converter separates dynamic range and resolution. This is well suited for audio applications where masking effects take place.

A floating point converter has been implemented in a low voltage 2 µm CMOS technology. Its power consumption is 50 µW at ± 1.25 V power supply, 16 kHz sampling frequency and the die area is less then 2 mm^2. The dynamic range is about 80 dB and a very good audio quality is obtained.

Though, some improvements could be performed in order to reduce the internal noise of the adaptive gain as well as to decrease its sensibility to technological parameters .

From a system level point of view, an elegant solution should be found to deal with null input signals. The adaptation table and strategy should also be investigated. Some experiments using a modified adaptation table showed that the reserve bit could be removed. The implications could be either a 6 dB higher SNR or a one bit smaller coarse quantizer or, better, a simpler adaptive gain. The two last propositions would lead to a lower power consumption.

The floating point converter might also be appropriate for other applications such as instrumentation. The adaptation table could be targeted for the processed signals and stored on-chip.

The floating point A/D converter for audio applications was successfully implemented. A number of questions are nevertheless still open and deserve further studies.

References

[1] Schaub A., "Micro Power A/D Converter for Audio Signals", *Digital Signal Processing*, Vol. 2, pp 110-114, 1992

[2] Jayant N. S., "Adaptive Quantization with One Word Memory", *Bell Syst. Tech. J.*, pp 1119-1144, Sept. 1973

[3] Grisoni L., "Micro Power 'Relative Precision' 15-bit A/D Converter", *ICSPAT'95*, pp 420-424, Boston, MA, 1995

[4] M. R. Schroeder, B. S. Atal, J. L. Hall, “Optimizing Digital Speech Coders by Exploiting Masking Properties of the Human Ear”, *J. Acoust. Soc. Am. 66, Dec. 1979*, pp 1647-1652.

[5] Heubi A., "Micro Power 'Relative Precision' 14-bit Cyclic RSD A/D Converter", *ISLPED'96*, Monterey, CA, 1996

[6] Ginetti B., " A CMOS 13-bit Cyclic A/D Converter", *IEEE Jour. of Solid-State Cir.*, Vol. 27, No 7, pp 957-965, July 1992

[7] Degrauwe M. “design of Analog CMOS IC’s”, 7th semester course, Microtechnology Institute, University of Neuchâtel.