

Targeting ultra-low power consumption with non-uniform sampling and filtering

Laurent Fesquet, Gilles Sicard

TIMA
CNRS-Grenoble INP-UJF
Grenoble, France
{Laurent.Fesquet, Gilles.Sicard}@imag.fr

Brigitte Bidégaray-Fesquet

LJK
CNRS- UJF-Grenoble INP- UPMF
Grenoble, France
Brigitte.Bidegaray@imag.fr

Abstract — Today signal processing systems uniformly sample analog signals without taking advantage of their intrinsic properties. For instance, temperature, pressure, electrocardiograms, speech signals significantly vary only during short moments. The digitizing system does not take into account this specificity and furthermore is highly constrained by the Shannon theory which fixes the sampling frequency at least twice the input signal frequency bandwidth. It has been proved that Analog-to-digital Converters (ADCs) using a non equi-repartition in time of samples leads to interesting power savings compared to Nyquist ADCs. A new class of ADCs called A-ADCs (for Asynchronous ADCs) based on level-crossing sampling (which produces non-uniform samples in time) and asynchronous technology has been developed. This article will present a fully non-uniform filtering technique associated to such an ADC which is able to drastically reduce the power consumption.

I. INTRODUCTION

Reducing the power consumption of mobile systems - such as sensor networks and many others electronic devices - by one to two orders of magnitude is extremely challenging but will be very useful to increase the system autonomy and reduce the equipment size and weight. This is particularly important for medical equipments such as pacemakers or mobile ECG recording systems. In order to reach this goal, this paper proposes a solution – based on non-uniform sampling - applicable to FIR filtering which completely rethinks the signal processing theory and the associated system architectures.

Signal processing systems uniformly sample analog signals at Nyquist rate without taking advantage of their intrinsic properties; this is a fact. Due to the Shannon theory, which fixes the sampling frequency at least twice the input signal frequency bandwidth, the digitizing system is highly constrained. Nevertheless, for many signals such as biological or medical signals (and many others), it will be useful to take into account their specific properties in order to determine the appropriate sampling scheme. For instance, a level-crossing

sampling scheme will be more adapted to a sporadic signal than to a modulated RF signal. Indeed, the sporadic signal will produce a very limited sample number contrarily to the RF signal. This low number of samples can be beneficial for limiting the consumed power. It has been proved in [1] and [2] that Analog-to-digital Converters (ADCs) using a non equi-repartition in time of samples can lead to interesting power savings compared to Nyquist ADCs. A new class of ADCs called A-ADCs (for Asynchronous ADCs) based on level-crossing sampling [3][4] and related signal processing techniques [5][6] have been developed.

After presenting the level-crossing sampling techniques and the associated A-ADC, which has been designed and fabricated in STMicroelectronics 130 nm CMOS technology, this paper suggests an important change in designing FIR filters. Like analog signals which are usually sampled uniformly in time, the filter transfer function is also regularly sampled with a constant frequency step. Nevertheless, non-uniformly sampling this transfer function leads to an important decrease of the weight-function coefficient numbers. Combined with a non-uniform level-crossing sampling technique performed by an A-ADC, this approach drastically reduces the computation load by minimizing the number of samples and operations.

II. ASYNCHRONOUS ANALOG TO DIGITAL CONVERTERS

This new class of converters named A-ADCs is based on two essential principles: a non-uniform sampling scheme of the continuous time signal and an asynchronous implementation.

A. Non uniform sampling

The principle of regular sampling is presented in Fig. 1a: samples are equi-spaced in time because sampling is ordered by an external clock of a fixed period T_{sample} . For non-uniform sampling (cf. Fig. 1b), 2^M-1 quantization levels are regularly set along the amplitude range of the signal (M defines the

hardware resolution of the converter). A sample is captured only when the analog input signal V_{in} crosses one of these levels. J.W. Mark *et al.* named this principle “level crossing sampling scheme” [1]. Contrary to the classical Nyquist sampling, samples are not regularly spaced out in time, because it depends on the signal variations: the sharper the signal, the closer the samples. Thus, together with the value of the sample b_i , the time Dt_i elapsed since the previous sample b_{i-1} must also be recorded, according to the resolution T_C of a time basis.




Fig. 1: Regular sampling (a) vs. non uniform sampling (b).

Uniform sampling (cf. Fig. 1a) is always implemented in classical A/D conversion systems: samples are taken by the Sample and Hold (S/H), respecting the Shannon theorem, then quantized by the ADC. Considering an ideal clock and an ideal S/H, time instants are perfectly known, the only imprecision is due to the quantization noise added during the A/D operation. It is characterized by the Signal-to-Noise Ratio (SNR), which only depends on the resolution N of the converter. The classical formula of the SNR for the Nyquist sampling scheme is: $SNR_{dB} = 10 \cdot \log(12 \cdot \sigma_{in}^2 / q_{Nyq}^2)$, where σ_{in} is the standard deviation of the input signal V_{in} and q_{Nyq} is the quantum. For a non-uniform sampling A/D Converter, the conversion of samples is triggered when a reference level is crossed by the signal. The amplitude of the sample is then precise, but the time elapsed since the previous sample is quantized according to the time resolution T_C . Dt_i is known with an error δt , which belongs to the interval $[0, T_C]$. The SNR relation must be re-determined. First, the error δt in time can be translated into an error in amplitude δV according to: $\delta V = (dV_{in}/dt)\delta t$, where dV_{in}/dt is the input signal slope. Then dV_{in}/dt and δt can be considered as independent random processes, thus this quantization noise power becomes:

$$P(\delta V) = P\left(\frac{dV_{in}}{dt}\right) \cdot P(\delta t) \quad (1)$$

δt is a random variable uniformly distributed across $[0, T_C]$, thus:

$$P(\delta t) = \frac{T_C^2}{3} \quad (2)$$

Like in the synchronous case, the SNR is always defined as:

$$SNR_{dB} = 10 \cdot \log\left(\frac{P(V_{in})}{P(\delta V)}\right) \quad (3)$$

Using (1) and (2), we get:

$$SNR_{dB} = 10 \cdot \log\left(\frac{3 \cdot P(V_{in})}{P\left(\frac{dV_{in}}{dt}\right)}\right) + 20 \cdot \log\left(\frac{1}{T_C}\right) \quad (4)$$

The first term of (4) is only determined by the statistical properties of the input signal V_{in} . The SNR depends on the timer period T_C , and not on the number of quantization levels. Thus, for a given implementation of the non-uniform sampling A/D converter (a fixed number of quantization levels: $2^M - 1$), the SNR can be externally tuned by only changing the period T_C of the timer. For example, if the SNR has to be improved by 6,02dB i.e. must have one more bit of effective resolution (or ENOB – Equivalent Number Of Bit), T_C has to be divided by 2. Equation (4) can be used to determine the theoretical SNR of a pure sine wave, a speech signal, an ECG signal... The main differences of regular vs. non-uniform sampling are summarized in TABLE I.

TABLE I. CHARACTERISTICS OF BOTH TYPES OF SAMPLING.

	Regular sampling	Non-uniform sampling
Conversion trigger	clock	level crossing
Amplitude	quantized	exact value
Time	exact value	quantized
SNR dependency	number of bit	timer period
Converter output	amplitude	(amplitude, time)

Theoretically for the non-uniform sampling, the SNR can be improved as far as it is needed, by reducing T_C . The limit is in fact the accuracy of the analog blocks: they determine the precision of the quantization levels position in Fig. 1b. If these levels are known with an incertitude δv , this error must be added in the quantization noise in (1), and the SNR is therefore degraded. If an ENOB-bit resolution is targeted for the asynchronous converter, each analog block must be designed to reach the ENOB-bit accuracy when used in the circuit.

B. Asynchronous conversion

Many hardware strategies can be used to implement the level crossing sampling scheme described in the previous paragraph. The chosen one is a tracking loop enslaved on the analog signal to convert. For the crossing detection, the instantaneous comparison of the input analog signal V_{in} is restricted to the two nearest levels: the one just above, and the one just below. Every time a level is crossed, the two comparison levels are down or up shifted. The conversion loop, shown in Fig. 2, is composed of four blocks: a difference quantificator, a state variable (an up/down counter), a Digital-to-Analog Converter (DAC), and a timer delivering the time intervals Dt_i between samples. This choice is very interesting

for a minimization of the hardware. Moreover, only one cycle of the loop is needed to convert a sample. As it is explained before, no external signal as a clock is used to trigger the conversion of samples. To preserve the same state of mind, an asynchronous structure has been chosen for the circuit. The information transfer between each block is locally managed with a bi-directional control signalling: a request and an acknowledgement. This explains the name of Asynchronous ADC. This architecture has been implemented in STMicroelectronics 130 nm CMOS technology.

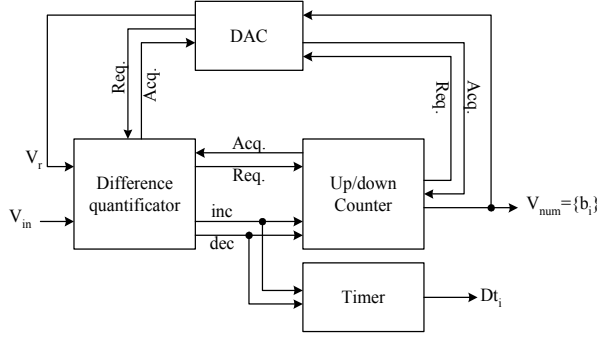


Fig. 2: Block diagram of the A-ADC.

III. FILTERING PRINCIPLES

As mentioned above, non-uniform sampling leads to a reduced number of samples, compared to the classical Nyquist sampling. This feature has already been used in [5][6] to design FIR filters. These techniques are based on interpolation and/or resampling. In these works, the authors however used a classical filter that follows the usual uniform discretization in frequency of the impulse response. Here we want to go further and take advantage of the fact that the filter transfer function (the Fourier transform of the impulse response) is a very smooth function with respect to frequency. Indeed, it can therefore be well approximated by a linear interpolation of quite few samples.

The signal which we want to filter is given in the time domain and is denoted by $x(t)$. The filter transfer function is given in the frequency domain and is denoted by $H(\omega)$. The result of the filtering process $y(t)$ is then theoretically the convolution of $x(t)$ with the impulse response $h(t)$ which is the inverse Fourier transform of $H(\omega)$:

$$y(t) = \int_{-\infty}^{+\infty} h(t - \tau)x(\tau)d\tau \quad (5)$$

The signal $x(t)$ is sampled using a level crossing scheme in the time domain. This scheme is also applied to the filter transfer function $H(\omega)$ in the frequency domain. As the transfer function is complex valued, we can choose to sample either when the amplitude crosses some predefined values, or the phase, or both. The samples read $(s_n, \delta t_n)$ for the signal and $(H_k, \delta \omega_k)$ for the filter transfer function. These samples are formed of a value and the (time or frequency) interval length "elapsed" since the last sample. To give results or describe algorithms we will use the sample times or frequencies defined as:

$$t_n = t_0 + \sum_{i=1}^n \delta t_i \quad (6)$$

$$\omega_k = \omega_0 + \sum_{j=1}^k \delta \omega_j \quad (7)$$

Notice that the computations will be performed only using the time and frequency intervals δt_n and $\delta \omega_k$. As mentioned above, the transfer function – thanks to its smoothness – can be well approximated by a piecewise linear function $\hat{H}(\omega)$. The corresponding impulse response $\hat{h}(t)$ can easily be split in several contributions $h_k(t)$ corresponding to each frequency sample:

$$\hat{h}(t) = \sum_k h_k(t) \quad (8)$$

This leads to a summation formula [7], as in the classical FIR filtering case, which takes the form of a convolution product.

$$\hat{y}(t) = \sum_n x_n \sum_k h_{nk}(t) \quad (9)$$

The effective formula directly depends on the type of interpolation which has been used (piecewise constant, linear, spline ...). In the sequel, we will use a piecewise constant interpolation.

IV. A DEMONSTRATIVE EXAMPLE!

In order to demonstrate the relevance of the approach, we choose a particularly simple formulation: the ideal low-pass filter which is 1 on the frequency interval $[-\omega_c, \omega_c]$ and zero elsewhere. Indeed, this yields a single sample $(1, 2\omega_c)$! This case is simple due to its minimal number of samples in the frequency domain, but it displays all the difficulties of the general case, i.e. the need to evaluate special functions. The computation of the filter coefficients and the determination of these special functions is not addressed in this paper but is presented in [7]. These functions are built in many libraries in view of a numerical implementation. This feature makes possible the construction of efficient lookup tables for a hardware implementation. We also can notice that this approach is able to provide an implementation of an ideal filter.

V. NUMERICAL RESULTS

To illustrate this simple example, we filter the signal $x(t) = 0.45 \sin(2\pi t) + 0.45 \sin(10\pi t) + 0.9$ with our ideal low pass filter with the cutoff frequency $\omega_c = 4\pi$ rad/s. The theoretical result is therefore supposed to be $y(t) = 0.45 \sin(2\pi t) + 0.9$. Even if this signal is not a typically sporadic (more samples are taken), we perform the computations within the Matlab SPASS (Signal Processing for ASynchronous Systems) framework [8]. This signal is sampled with a 3-bit Asynchronous A/D Converter (AADC)

which leads to a level crossing sampling over the amplitude range [0V, 1.8V].

We can choose as we want the times at which the filtered signal is computed. To display the results we choose the sequence of times $t_m = .17m$ where m is an integer to have sampling points dispatched irregularly over the obtained solution. On Figure 3, you can see the result for a linear interpolation of the signal non-uniform samples and a 3-bit AADC. We plot continuous functions with lines: the initial signal $x(t)$ (dashed line) and the theoretical filtered signal $y(t)$ (solid line). We plot the sampled results with markers: the non-uniformly sampled initial signal x_n (asterisk markers) and the computed filtered samples y_m (circle markers) at times t_m .




Fig 3: Initial signal (dashed line), theoretical filtered signal (solid line), non-uniformly sampled initial signal (asterisk markers) and computed filtered samples (circle markers)

This very simple test case has quite a low number of parameters compared to the full problem for which we can finely tune the filter transfer function sampling for example. We compare here the results obtained for a zeroth and a first order interpolation of the signal and for different values (2, 3, 4 and 5) of the AADC resolution. On Table II, we give the relative error between the computed filtered samples y_m at times $t_m = .01m$ (m integer) and the theoretical values $y(t_m)$.

TABLE II. Error of the filtering method for the 0th and first order interpolation of the signal and different resolutions (M) of the AADC.

	0 th order	1 st order
M = 2	6.08%	5.84%
M = 3	0.76%	0.48%
M = 4	0.52%	0.45%
M = 5	0.46%	0.45%

In the case of the 2-bit AADC, there are 2.8 points per wavelength for the highest frequency part of the signal. This is a very low rate, and we are however able to have only 6% error on the filtered result which is quite sufficient for a large range of applications. With 3-bit A-ADC, the results show less

than 1% of relative error. This is usually enough accurate for most of the mobile and autonomous applications requiring a digital filtering.

VI. CONCLUSION

This article has presented the principles of the level crossing sampling scheme and an architecture of A-ADC able to implement it. This technique is especially well-suited for processing sporadic signal such as many medical or biological signals (ECG, EEC ...). Indeed, these signals are favorable to drastically reduce the number of samples and thus the global power consumption for processing these signals. Then we have presented a novel approach to FIR filtering based on the non-uniform sampling (in time) of the signal and the non-uniform sampling (in frequency) of the filter transfer function. This non-conventional sampling of the transfer function leads to an important reduction of the filter coefficient number. Even if the computation of these coefficients is more complex, the filtering process only requires very few operations for calculating the output signal. This low computation load results in a drastic power saving compared to the classical signal processing approach.

Previous works have demonstrated the feasibility of an ADC based on a level sampling scheme. The theoretical study for developing new filtering techniques has been evaluated using Matlab simulations. We now are currently working on a hardware implementation. This approach is very promising to achieve ultra-low power consumption in mobile and autonomous systems, especially with biological signals which often display a sporadic nature.

ACKNOWLEDGMENT

This work has been supported by a funding from the Joseph Fourier University, Grenoble, France: project TATIE.

REFERENCES

- [1] J. W. Mark and T. D. Todd, "A nonuniform sampling approach to data compression," IEEE Trans. Commun., vol. 29, pp. 24–32, January 1981.
- [2] N. Sayiner, H. V. Sorensen, and T. R. Viswanathan, "A level-crossing sampling scheme for A/D conversion," IEEE Trans. Circuits Syst. II, vol. 43, pp. 335–339, April 1996.
- [3] E. Allier, G. Sicard, L. Fesquet, and M. Renaudin, "A new class of asynchronous A/D converters based on time quantization," in Proc. ASYNC'03, Vancouver, Canada, May 2003, pp. 196–205.
- [4] F. Akopyan, R. Manohar, and A. B. Apsel, "A level-crossing flash asynchronous analog-to-digital converter," in Proc. ASYNC'06, Grenoble, France, March 2006, pp. 11–22.
- [5] F. Aeschlimann, E. Allier, L. Fequet, M. Renaudin, "Asynchronous FIR Filters: Towards a New Digital Processing Chain", Proceedings of the ASYNC Conference, pp. 198-206, April 19-23 2004, Crete, Greece.
- [6] F. Aeschlimann, E. Allier, L. Fesquet, and M. Renaudin, "Asynchronous FIR filters: towards a new digital processing chain," in Proc. ASYNC '04, Crete, Greece, April 2004, pp. 198–206.
- [7] B. Bidégaray-Fesquet and L. Fesquet, "A fully non-uniform approach to FIR filtering," Sampling Theory and Applications Conference, Marseille, France, May 2009, 4 pages.
- [8] <http://ljk.imag.fr/membres/Brigitte.Bidegaray/SPASS>