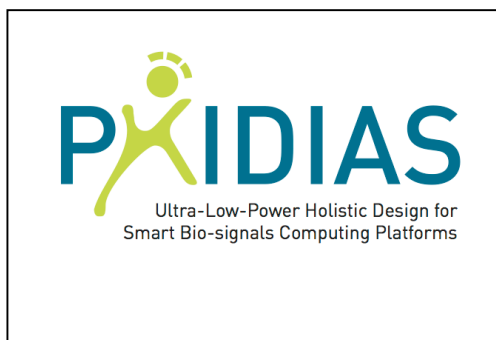Deliverable report for

P H I D I A S

"Ultra-Low-Power Holistic Design for Smart Biosignals Computing Platforms"

Grant Agreement Number 318013

Deliverable D 2.1
Report on feasibility of analog compressive sampling platform implementation for consortium use
Due date of deliverable: 30/06/2013

Lead beneficiary for this deliverable: IMEC-NL
Contributors:
- **IMEC-NL:** analog front-end model and investigation of energy efficiency in different scenarios, low-power memories
- **UNIBO:** interconnect topology and interconnect implementation
- **EPFL:** digital signal processing core

Dissemination Level:		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Version: 1.3 (revision)	Date 04.03.2014
Draft of the WP Leader	
Commented version for amendment	
Version accepted by the Steering Board	X

1. Description of task

This report describes the initial feasibility study of an Analog Compressive Sampling Platform (ACSP) for biopotential signals. A model for the estimation of the power consumption has been developed which captures the essential trade-offs between the key parameters such as the noise, bandwidth, resolution and the power consumption. Comments about the other parameters such as the area requirements and impact of technology scaling on implementation of ACSP have been added at a higher level and the detailed analysis of those aspects are left out for future study. Brief description of the rest of the architectural components such as Digital Signal Processor (DSP), interconnect etc. have also been added for sake of completeness in line with tasks defined in T2.1.

2. Description of work & main achievements

2.1. Compressive Sampling Platform

The Phidias project aims to investigate the opportunities offered by exploiting sparsity of bio-signals in order to efficiently acquire and compress them. Compressive sampling (CS) is one of the techniques which uses the sparsity of the signals in a particular basis to achieve this objective.

CS algorithm, in general, can be implemented either in analog or digital domain. The choice of implementation (in analog or digital) is determined by various factors such as the power consumption, area, technology, and others. It must be noted that each of the implementations have their limitations arising from the fundamental limits. For example, analog implementations are generally limited by device noise and non-linear characteristics and in deep submicron technology are limited by mismatch and the voltage headroom available. Digital circuits on the other hand are robust to noise but their power consumption increases with the data-rate. In deep submicron technology the power consumption of the digital circuits is dominated by the leakage currents. In the current work, we focus on deriving a model for power consumption for an analog compressive sampling platform.

In deriving the models, we are required to make several assumptions and we tried to explicitly mention all the assumptions made and the rationale behind them, as far as possible. Extremely accurate models are usually difficult to come up with due to the large number of variables involved and a specific application has not yet been determined.
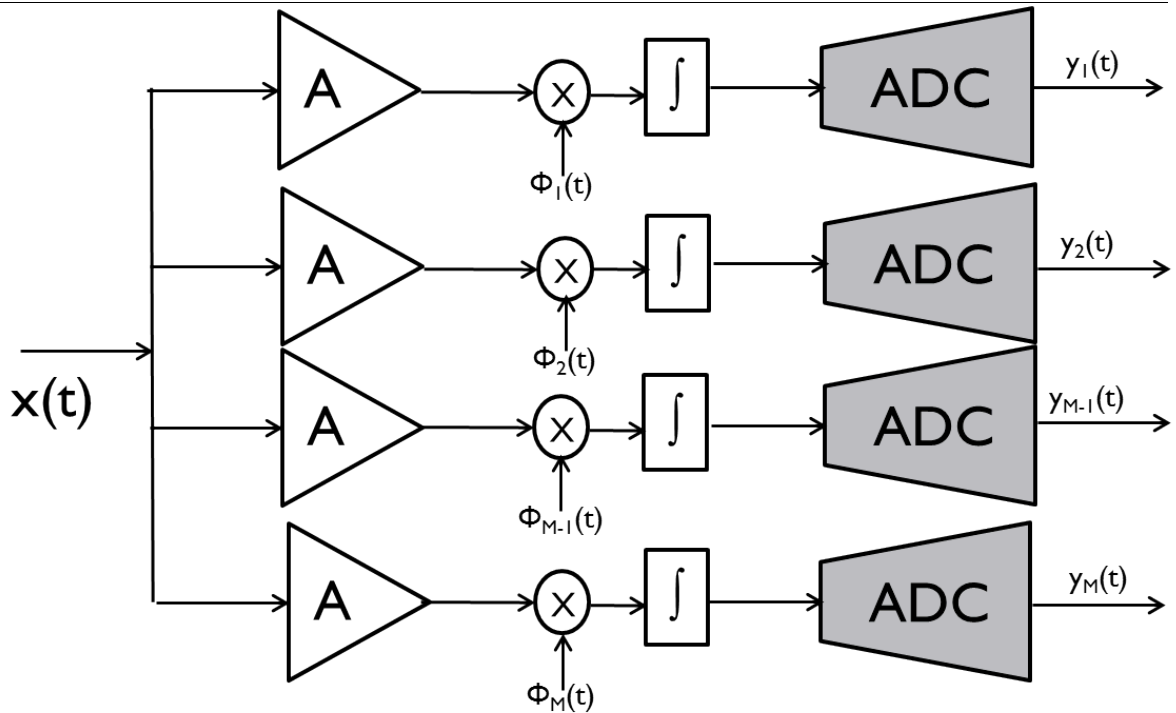
Figure 1: Analog Compressive Sampling Platform

2.2. Initial investigation of the power model of analog compressive sampling platform

Compressed Sensing (CS) exploits the sparse nature of biological signals for sampling at rates much lower than the traditional limits defined by the Nyquist - Shannon theory. In other words, CS assumes that bio-signals admit a very compact representation in a carefully designed base (or dictionary). In this Section, we provide an initial model for the estimation of the components involved in the ACSP and is to be considered as a baseline over which implementation specific details would be incorporated and evaluated.

2.2.1. Analog to Digital Conversion

Analog to digital converters (ADC) are indispensable components of the mixed signal systems, forming the link between the analog and digital systems. A typical ADC is characterized by the resolution of the converter and the bandwidth among several other parameters. Several architectures for the ADCs exist, which usually trade-off resolution for bandwidth and power consumption. Figure 2 shows the trade-off between the resolution and the bandwidth for various ADC topologies [1].

In order to compare the performance of a ADC across various topologies, a metric known as the *Figure of Merit* (FoM) was introduced which captures the trade-off involved between the resolution, bandwidth and power consumption, which is defined as

$$FoM = \frac{P_{diss}}{f_{samp} \times 2^B}$$

Where

P_{diss} = Power dissipated by the ADC

f_{samp} = Sampling frequency (For oversampling converters it is twice the bandwidth of signal)

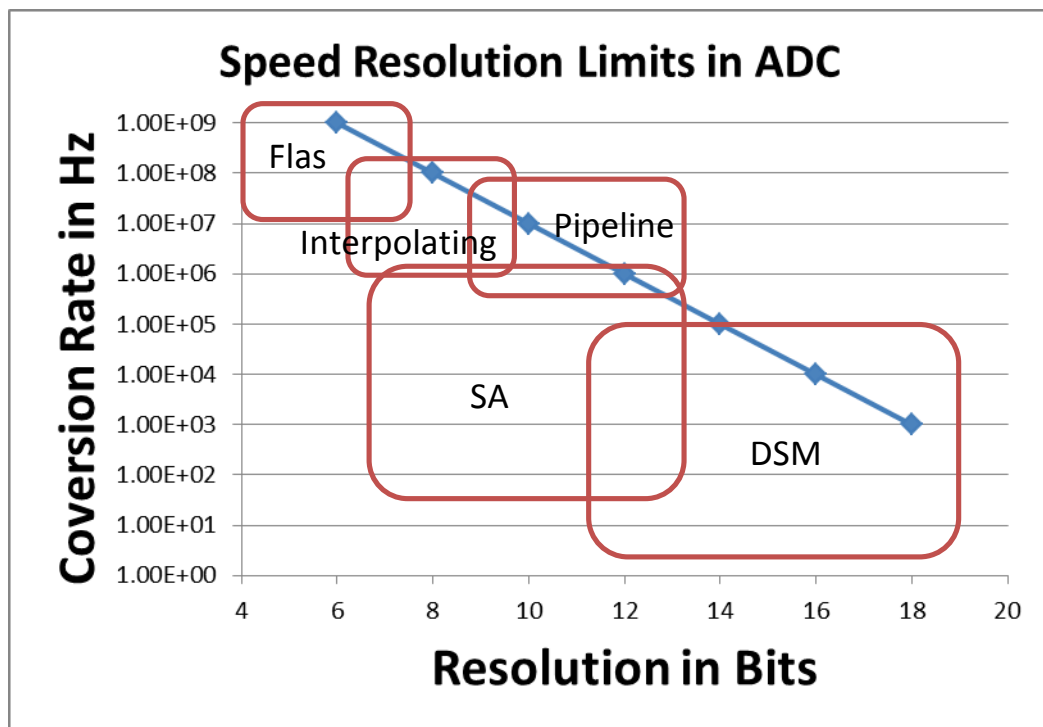B = Effective number of bits of the converter

Figure 2: Speed Resolution Limits in ADC

As evident from the definition of FoM, the power dissipated by a ADC doubles for every increase in one bit of the converter. This conclusion, however is true when the resolution is limited by quantization noise. For very high resolution converters, where the resolution is limited by thermal noise, the power of the ADC increases by a factor of 4 for every one bit increase in the resolution of ADC. However for most of the biopotential acquisition systems, the resolution of the converters employed are in the range of 10-16 bits where the resolution is quantization noise limited.

Typical ADC architectures employed in biopotential acquisition systems are Successive Approximation (SAR) and Delta Sigma Modulation(DSM) ADCs. SAR ADCs are used in moderate resolution and bandwidth applications while DSM ADCs are used in high resolution and low bandwidth applications. SAR ADCs generally tend to have a lower FoM (and hence lower power consumption at a given resolution and bandwidth) compared to a DSM ADCs due to the fact that the SAR architecture is digital intensive.

For an ACSP shown in Figure 1, the effective bandwidth of the signal seen by each

ADC is reduced by a factor of N, where N is the window size over which the random modulated signal is integrated. However, the number of ADCs in an ACSP is M times larger than a conventional system and hence the total power consumed by the ADCs in an ACSP system is given by

$$P_{adc} = \frac{M \times FoM \times 2^{Bm} \times f_{samp}}{N}$$

Where the terms mean their usual meaning. Note that the effective resolution of the converter now corresponds to the required resolution of the measurements (B_m) and *not* that of the resolution at which the biopotential signal is intended to be captured (B).

Therefore, we can define the Compression Ratio (CR) as follows.

$$CR = \frac{M \times B_m}{N \times B}$$

The choice M and N depend on the sparsity of the signal on the selected basis, desired reconstruction quality and the random modulation scheme that is chosen. From D.1.1 and studies conducted elsewhere, it is observed that a larger value of N gives better compression at a given reconstruction quality. However a larger value of N has certain problems, which are discussed under the Section 2.2.2.

Typical ADC technology trends show an increase in ADC efficiency by a factor of 2 for every two years [2]. In 2008, an ADC with a FoM of 100fJ/conv-step was considered state of the art. Going by the trend one can expect the current state of the art FoM for ADCs to be around 10fJ/conv-step (and it is indeed the case. In fact a SAR ADC with a FoM of 2.2 fJ/conv-step for up to 8-10 bit resolution has been demonstrated recently [3]). Henceforth for the estimation of the power consumed by the ADC, a FoM of 10fJ/conv-step is deemed to be appropriate.

Additional remarks of ADC power consumption

Generally, the power consumption that is often quoted for a ADC is the power consumption of the core components that make up the ADC such as the preamplifiers, latches and integrators etc. However almost all ADCs require references that are stable across variations in temperatures, power supply voltages, etc. (Bandgap references are often the most preferred). All the synchronous (and some classes of asynchronous) ADCs require a clock signal that is very stable (The jitter of the clock can reduce the Effective Number Of Bits (ENOB) of an ADC. The exact details of the influence of the jitter of ENOBs is specific to the architecture of the ADC). Moreover the ADC requires to be driven by the analog front end and usually a buffer is employed for the same. The FoM of the ADC does not include the power consumption of the *other* components that have been just mentioned. Henceforth is reality the ADC power consumption is several times higher than the value that is obtained from the FoM. The power consumed by these blocks can be as high as a few microwatts to tens of microwatts depending on the value and accuracy requirements. For typical biopotential applications, the power consumed by references, clock generators and buffer is of the order of 1-10 µW [4].

In an ACSP employing *M* parallel paths for signal acquisition as in Figure 1, it can be

expected that the system can work with a single reference (without loading it significantly), a buffer for each ADC and a single clock. However the clock generator has to now drive *M* ADCs and will consume more power compared to driving a single ADC.

2.2.2.   Front End Amplifier

The Front End Amplifier (FEA) in any system, forms the interface between the sensor and the rest of the system. Hence the performance of the overall system is a critical function of the performance of the FEA. In fact several of the vital parameters for the biopotential acquisition systems such as the noise, common mode rejection ratio (CMRR), power supply rejection ratio (PSRR) are set by the FEA.

In general, the biopotential signals are extremely small in magnitude. Typical values for Electro Cardio Gram (ECG) are 1-10mV and for Electro Encephalo Gram (EEG) are 1-100µV. This would imply that the input referred noise has to be extremely low to keep the signal above the noise floor. As mentioned earlier, the design of the FEA dictates the overall noise of the system.

In order to quantify the trade-offs between the noise, bandwidth and the power consumption of the amplifier, a figure of merit known as the Noise Efficiency Factor (NEF) is defined as follows [5].

$$ NEF = V_{rti,rms} \sqrt{\frac{2 \times I_{total}}{\pi \times U_T \times 4 \times kT \times BW}} $$

Where

$V_{rti,rms}$ = Input referred rms value of the noise voltage

I_{total} = Total current consumption of the amplifier

T = Absolute temperature

k = Boltzman constant

U_T = Thermal voltage defined as kT/e where e is charge of the electron

BW = Bandwidth of the amplifier

As evident, a minimum value of NEF indicates a power efficient amplifier for a given specification of noise and bandwidth. The minimum achievable NEF factor for a given topology depends on the subthreshold slope of the MOSFETs used to make the amplifier and thereby limited by technology. Typical values for the NEF for various biopotential amplifiers range between 2 and 5. More recently alternate architectures which are termed 'current reuse structures', 'feed-forward noise cancellation' and 'signal nulled noise feedback' have achieved a NEF of the range 1.5-2 [6]. Therefore a value of 2 for the NEF is suitable for the estimation of the power of amplifier.

It must be noted that the NEF considers the total current consumption as the

representative for the power consumption of the amplifier. However the power consumption is also a function of the power supply voltage (as shown below). Hence for the amplifiers designed in deep submicron technology, an alternate metric to capture the trade-offs involved has been proposed as $NEF^2 \times V_{DD}$ [7], where NEF is same as the one defined earlier and V_{DD} is the supply voltage. For the current study, we discard the effects of technology and voltage scaling and reserve it for future study.

Therefore the power of the amplifier for a give input referred noise and bandwidth is

$$P_{amp} = NEF^2 \times \frac{V_{DD}}{2 \times V^2_{rti,rms}} \times \pi \times U_T \times 4 \times kT \times BW$$

To achieve an effective resolution of B bits, the input referred noise of the amplifier and the signal processing chain has to be less than the quantization noise of the ADC referred to the input. Assuming a full scale swing at the input of the ADC, the quantization step size at the ADC can be computed as

$$\Delta = \frac{V_{DD}}{2^B}$$

And the mean square value of the quantization noise voltage is given by

$$V_{QN,ms} = \frac{\Delta^2}{12}$$

Let G be the voltage gain of the amplifier and signal processing chain combined until the input of the ADC. Then the mean square value of the quantization noise referred to the input is given by

$$V_{QN,rti,ms} = \frac{V_{QN,ms}}{G^2}$$

Typical values for the value of G for biopotential acquisition systems range from 100-500 (V/V) and even 1000 depending on the supply voltage and the strength of the input signal (V_{in}). Usually the value of gain is selected such that the input of the ADC can swing from rail to rail. This condition can be expressed as

$$V_{DD} \geq G \times V_{in}$$

However in an ACSP, it is not the signal that will be quantized but the random modulated and integrated value of the signal over a window of N samples. Assuming that the samples after random modulation are independent and identically distributed, one may conclude that the instantaneous magnitude of the signal to be quantized is $\sqrt{N}$ times larger on an average. Hence the constraint that the voltage swing at the ADC has to be within supply voltage can be rewritten as

$$V_{DD} \geq \sqrt{N} \times G \times V_{in}$$

Usually the larger value of N facilitates higher compression (as demonstrated in

D.1.1). But this restricts the maximum possible gain that AFE can provide. Compromise on the value of gain can result in higher input referred noise and therefore a trade-off exists between the compression ratio and noise given the magnitude of input signal and supply voltage.

It must be emphasized that it is not necessary that all the blocks in signal processing chain to add a gain. For example, some of the mixers do attenuate the signal and even exhibit a time varying gain (refer to the section on mixers for details). For the analysis we only consider the average value of the gain.

The constraint on the input referred noise with respect to the quantization noise can be expressed as

$$V^2_{rti,rms} \le \frac{V_{QN,ms}}{G^2} = \frac{V^2_{DD}}{2^{2B} \times 12 \times G^2}$$

Additional remarks

Apart from the bandwidth, noise and power consumption there are several other characteristic parameters of an amplifier that are very relevant for biopotential acquisition systems. The input impedance of the amplifier is one of them. Usually it is desirable to have a very large input impedance for the amplifier to avoid loading of the electrodes and also to mitigate the effects of mismatch in the impedance of electrodes as follows. Consider the equivalent circuit of an amplifier with its input impedance as shown in Figure 3. Let Ze1 and Ze2 be the impedance of the electrode 1 and 2 respectively and Zin be the input impedance of the amplifier.

Ideally, the amplifier is supposed to amplify the difference between the input voltages ie (v1-v2) and reject the common mode voltage completely. However, in the presence of the electrode mismatches, the common mode component of the signal
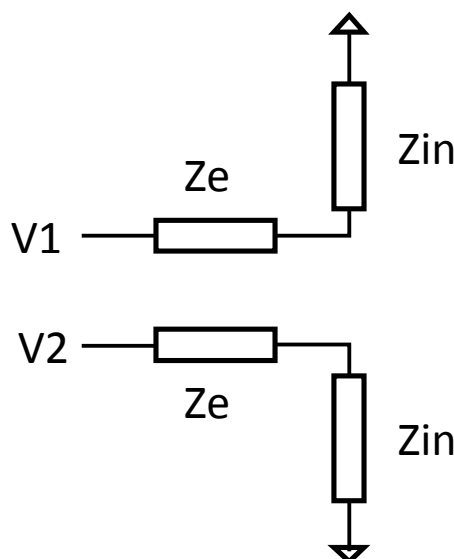
Figure 3: Effect of Electrode Mismatch on Amplifier Common Mode Rejection

appears as differential component as described below.
Let V1 and V2 be the potentials sensed by electrodes 1 and 2 respectively with

respect to a common reference. Then the differential voltage sensed by the amplifier in presence of electrode impedance mismatch is given by

$$V_d = V_1 \frac{Zin}{Ze1 + Zin} - V_2 \frac{Zin}{Ze2 + Zin}$$

This implies that the common mode component of the input signal gets converted to the differential component and this effect can be minimized when *Zin >> Ze1, Ze2*. However in case of an ACSP with *M* parallel paths, the effective input impedance seen look into the input of the amplifier is reduced by a factor of *M*. This would mean that the ACSP has an inferior common mode rejection. The problem is more pronounced when the electrodes have a high impedance (as in the case of dry and polymer based electrodes). Impedance boosting techniques such as bootstrapping exist to increase the input impedance of the amplifiers significantly (of the orders of 100 GΩ) but they come at the expense of increased power consumption.

Another aspect, that is intrinsic to the MOSFETs (strictly speaking MOSFETs carrying a DC current), that is of concern to the biopotential acquisition systems is the flicker noise. Flicker noise has a power spectral density that has a 1/f characteristic and hence is predominant at low frequencies, which are typically the frequencies of several biopotential signals. Hence it is almost often the case in biopotential acquisition systems, care has to be taken to reduce the flicker noise. One of the several possibilities include use of pMOS devices for the input pair, making the devices larger and using techniques like 'chopping' and 'correlated double sampling' (CDS). Among the techniques presented, chopping is the most preferred due to its simplicity. However, chopping requires a clock signal (can be derived from the system clock) and power to drive the switches. Chopping also results in an increase of white noise floor (due to modulation of white noise around the harmonics of chopping frequency).

Other aspects such as the Electrode Offset Voltage (EOV) requires servo loops for cancellation which consume additional power that is not captured by NEF.

2.2.3. Mixer and Integrator

The mixer provides the pseudo-random modulation to the signal. This is performed by multiplying the amplified signal from the sensor interface (RF) with a random signal (LO). The random signal can be implemented with a square wave oscillating between ±1 with random transitions during a cycle. Although the ADCs sample at a rate effectively lower than the Nyquist frequency, the mixer (as well as the other analog front-end) must operate at or above the Nyquist rate.

The two primary mixer topologies are current commutating or voltage commutating. The main difference is that current commutating mixers have a g_m stage. One common current switching mixer is the Gilbert cell, and a common voltage switching mixer is the ring mixer. The diagrams for these topologies are shown in Figure 4.

A slight alternative to the passive ring mixer is to replace M1 and M2 with PMOS devices as in an H-bridge configuration. The gates of these transistors would also be connected to the opposite LO polarity as M1 and M2 shown in Figure 4.

The Gilbert cell is advantageous because it has positive conversion gain, it does not require a differential RF signal, and it offers superior isolation from RF feedthrough. However, it requires a noisy and large passive element, and suffers from nonlinearity. The passive mixer is advantageous because it has no DC current through the transistors and therefore has no flicker noise (although this is bias dependent). It also has very high linearity because it does not have a transconductance stage, and has a very small area. However, it requires a large amount of power to drive the switching devices, it is lossy (negative conversion gain), and requires a differential sensor signal which may be difficult to obtain in certain applications.
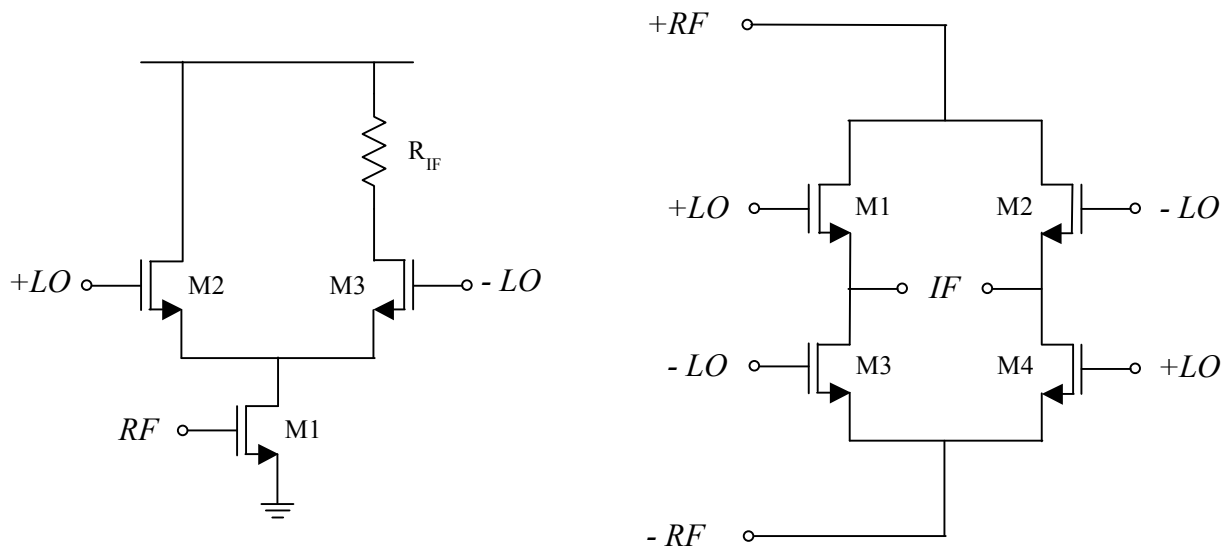

Figure 4: Gilbert cell (left) and passive ring mixer (right)

Conversion gain

All conversion gain calculations are made assuming that the LO signal is a 50% duty cycle square wave. This is not the case for a ACSP, but it is adequate for feasibility analysis.

For the single-ended Gilbert cell, the conversion gain (Gc) is simply the transconductance of M1 (gm1) multiplied by the factor $1/\pi$. If the output were taken differentially, it would be multiplied by $2/\pi$. It is expressed in the equation below

$$G_c = \frac{g_{m1}}{\pi}$$

For the passive mixer, it is necessary to create a Thevenin equivalent circuit. The equivalent voltage source ($V_T(t)$) and conductance ($g_T(t)$) are given by equations below. g(t) is the time-varying conductance of the switches M1-M4.

$$V_T(t) = \frac{g(t) - g(t - \frac{T_{LO}}{2})}{g(t) + g(t - \frac{T_{LO}}{2})} \times V_{RF}(t)$$

$$g_T(t) = \frac{g(t) + g(t - \frac{T_{LO}}{2})}{2}$$

With a square wave LO signal and zero capacitive load, the conversion gain G_c is $2/\pi$ or -3.8dB. It is possible to design a passive mixer gain, but the theoretical limit is 0dB. Because the ACSP mixer effectively has a time-varying frequency, the gain will also vary with time.

Power consumption

The only source of power consumption is the switching of the transistors M1-M4. However, it is necessary to switch these gates from rail to rail. For the purposes of this study, only the power consumption of the passive mixer will be analysed. Although it may be necessary to drive the gates of M1-M4 for high speed switching, the power of the mixer is simplified to the formula given below

$$P_{mixer} = 4 \times C \times V_{DD}^2 \times f_s \times \alpha$$

Where C is the capacitance of each of the switching transistors, VDD is the supply voltage, fs is the switching frequency, and α is the activity factor. For a typical mixer, the LO signal is a clock and the activity factor is 1. However, for a randomly modulated signal, the probability of a 0 to 1 transition in a given cycle is 0.25. The sizing of the mixer devices depends heavily on the integrator, which is implemented here as a sample and hold capacitor before the ADC.

Integrator

The value of integrating capacitance (C_{int}), and therefore the value of N (where N is the number of averaged Nyquist samples) drives the design of the mixer. A larger N requires a large capacitance, and therefore requires large mixer switches and a larger physical area. However, a small N will have higher kT/C noise from the amplifier and mixer and will have worse compression characteristics because of the law of large numbers.
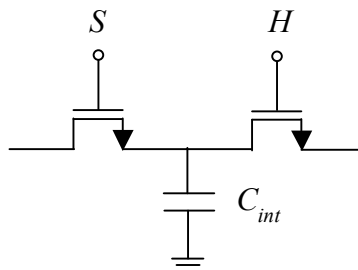
Figure 5: Sample and Hold Integrator Circuit

Another consideration to keep in mind is the amplifier gain with respect to the C_{int} capacitance. N samples accumulate on this capacitance, so if the gain is too high, the voltage will saturate. This can be solved by having different voltage rails for the ADC and the amplifier, by using a larger capacitance, or by using the same capacitance and reducing N.

2.3.   Area and Technology Considerations

Feature size scaling in CMOS technology has kept up with Moore's law over the past several decades. An immediate impact of scaling is the reduced voltage levels, which translate into a limited signal voltage swing, which translates into a reduced signal to noise ratio. This limits the achievable dynamic range for conventional voltage domain architectures at a given power.

Alternate architectures which use the representation of signal on alternate domains (such as current, charge, time and phase) have been proposed in the literature to alleviate the problem of supply voltage reduction. Using such techniques, analog front ends that operate from supply voltage as low as 0.5V have been demonstrated with dynamic ranges exceeding 60dB. However the NEF for such amplifiers is generally higher than the conventional amplifiers.

Usually it is believed that the digital circuits benefit from the technology scaling due to the reduction in the gate and other parasitic capacitances which needs to be charged/discharged (referred to as dynamic power). Although it is true that the dynamic power consumption of a digital circuit decreases with the scaling, the leakage power dominates at the deep submicron feature sizes. Several techniques exist in literature to mitigate this problem. However the point to be emphasized is that the choice of technology of implementation impacts the architectural considerations as well as the power consumption of the system.

From the Figure 1, it can be seen that the ACSP employs M parallel acquisition paths, each having an integrator and an ADC of its own. Typically in integrated circuit implementation capacitors dominate the chip area and capacitors are almost always used in integrators and ADCs. Hence it is reasonable to expect the ACSP to occupy a larger area compared to a digital implementation.

2.4.   Processor model

The investigated processor (TamaRISC, being developed by EPFL) will be optimized for ultra-low-power operations, employing deep sleep modes and aggressive voltage-frequency scaling. To further reduce energy consumption, the datapath width will be constrained to 16 bits, while the pipeline depth will be small (3 stages in the current implementation). Forwarding paths will be used to limit the number of stalls during execution.

Care is being taken to retain a compact instruction set, which comprises 18 instructions in the baseline implementation, in order to reduce complexity and area, with the goal of increasing run-time efficiency. In fact, the baseline TamaRISC core only requires 72.3 kilo-gates. Nonetheless, the benefit (in energy terms) of employing a small number of specialized instructions for executing Compressive Sampling will be explored in the MS1 report.

The architectural exploration effort will continue in the second year of the project, comprehensive results, from an area, timing and energy consumption viewpoint, will be described in D2.3, due on month 24.

2.5.   Memory design and modelling

Design of efficient memory subsystem at architectural level is hampered with the lack of accurate models of memory blocks. Such a design becomes even more challenging at scaled technology nodes when standby power, PVT variability, and reliability are to be considered early on in the design stage in order to guarantee an energy-efficient design.

Further difficulty in memory design arises in the context of systems with extreme power constraint such as the one used in wireless sensor nodes (WSN) with limited available power source. An emerging approach of reducing the power in WSN is to operate the design in very low voltage close to the threshold voltage of the transistors. Commercial SRAM bit cells fail at such a low voltage making the operation of complete system at a single low voltage supply not possible.

In Phidias, we address two main challenges mentioned above, namely:

Design of memory modules capable of operating at near threshold voltage

An accurate memory models for the designed memory modules

In the first phase of Phidias we focused on the former challenge; We had a fully synthesizable memory module based on standard cell modules. An instance of this design (a 1Kx16bit bank) was taped out in Q4 2012. On the same tape out, a commercial 6-T memory of identical size was also put for comparison. A built-in-self test was implemented to enable testing the memory without being constrained by the operating voltage of IO cells. We have done the silicon measurements comparing the two memories minimum operating voltage. The second phase, we will develop a model of the proposed memory calibrated with measurement and simulation results. An accurate modelling of different performance metrics (e.g. delay, power, and area) during design space exploration is essential to make the right design decisions. The memory models will be used in WP3, where overall system design is done.

2.6.   Logarithmic Interconnect model

The goal of the logarithmic interconnect (being developed by UNIBO) is to offer a high-bandwidth, low-latency interconnection to enable fast communication between the processing elements and the multi-banked memory in the ULP architecture. To enable such behaviour the devised architecture for the logarithmic interconnect consists of a Mesh-of-Tree (MoT): a routing stage determine the routing path (based on address decoding) and an arbitration stage will handle, in a round-robin fashion, conflicting requests. Such architectural template offers benefits for Compressed Sensing execution, where different cores process data from different leads and a careful static allocation of data can deploy the interconnect characteristics to greatly increase the throughput.

To enable a single-latency communication between cores and the memory, the initial design explorations devise a fully combinatorial network. Paths between cores and memory are likely to be among the critical paths of the design and, in case of variations, can lead to timing violations. Since near-threshold operation drastically increases the sensitivity to variations, UNIBO is putting research effort in devising

architectural solutions to increase the resiliency of the ULP architecture. The idea is to have a runtime mechanism capable of dynamically reconfigure the latencies so as to tolerate variations.

The design and exploration efforts will continue in the second year of the PHIDIAS project and final results will be presented in deliverable 2.4, due on month 30.

3. Performance of the partners

All partners fulfilled their tasks in satisfactory time and quality. The contributions from EPFL on processor model and UNIBO's contribution on interconnect are acknowledged.

4. Conclusions

It is expected that the most of the power consumed by the ACSP is in the amplifiers and hence the design of energy efficient amplifiers is key to reducing the overall power of the system. Techniques that result in a lower value of NEF are suggested to be employed in amplifier design and study the impact of power reduction and other trade-offs invloved. Techniques to couple the output of a single transconductance stage to multiple transimpedance stages can also be explored.

To have a good trade-off between the power dissipation and bandwidth at moderate resolutions (10-12 bits) SAR ADC architecture is suggested to be used for the ACSP.

Given that the nonlinearity of the mixer is the most detrimental non-ideality to system performance, a passive mixer is recommended in spite of the lack of gain and potentially higher power consumption. If a technique to modulate the LO at sub-Nyquist frequencies exists (such as the spread spectrum random modulation), it can further reduce the power consumption of the mixer.

One major trade-off in the system design is the choice of N. A large value will have higher power consumption and will require a large S/H capacitance in the integrator. However, it will also improve the compression ratio because a larger random matrix will have improved incoherence.

The sizing of the mixer and integrator also presents design trade-offs. Larger S/H transistors will have higher bandwidth at a given thermal noise level and hence lower settling error. However, they will have greater charge injection, higher clock feed-through, and will present a larger capacitive load to the mixer. Larger devices also need to be driven and have higher dynamic power dissipation.

In order to quantify the above statements and to estimate the approximate power consumption, two application scenarios have been studied as below. The study is confined to ECG signals.

Scenario 1:

This scenario concerns about the ECG acquisition system for medical grade applications. For medical grade ECG systems, ANSI AAMI standards define the system specifications [8]. The following specification is better than the minimum requirements specified by standards.

Parameter	Specification
Noise	1uV rms referred to input (RTI)
Bandwidth	256 Hz
Resolution	16 bits

For such high resolution applications, usually DSM ADCs are employed due to their ability to shape the quantization noise. A mentioned in Section 2.2.1, a DSM ADC has inferior FoM compared to a SAR ADC. Typical numbers of the FoM for a DSM ADC range from 35fJ/conv-step to 75fJ/conv-step. For the conservative estimate we use 50fJ/conv-step for the analysis.

Also given the high resolution requirements, it is imperative that the open loop gain of the FEA has to be large (>140 dB for a 40dB closed loop gain). This puts a limit on the architectures that could be employed for the FEA. Generally a two stage amplifier with cascoded first stage is used to get such high gain. Such architectures generally have a higher NEF, thanks to the fact that power has to be spent in moving the non-dominant pole to higher frequencies. A NEF in the range 3-5 is more appropriate for such architectures and for the current calculation a NEF of 4 is used.

Following the reasoning from the above and employing the equations derived in Chapter 2, the power consumption for a single path in the ACSP is computed as follows.

$$P_{adc} = \frac{FoM \times 2^{Bm} \times f_{samp}}{N}$$

Using a value of 18 for Bm and 256 for the N, the ADC power turns out to be 26nW per channel.

Similarly, the power of the FEA is computed as follows

$$P_{amp} = NEF^2 \times \frac{V_{DD}}{2 \times V^2_{rti,rms}} \times \pi \times U_T \times 4 \times kT \times BW$$

Since a DSM ADC is employed, it can be safely approximated that the in-band input referred quantization noise is negligible compared to the thermal noise and henceforth we approximate that the limit on the resolution is dictated by the thermal noise. For 1.8V supply, with a gain of 100, the maximum signal swing that can be allowed at the input is 18mV peak to peak. This sets a limit on the thermal noise floor to 75nV rms RTI for a 16 bit resolution requirement. Using these numbers the power required for FEA turns out to be 825uW per channel.

Due to the dynamic nature of the mixer, in order to estimate the power accurately,

simulations have been performed with TSMC 0.18u devices with a V_{DD} of 1.8V and the power of the mixer turns out to be 3pW per channel, for a 1ux1u switch.

Summing the powers of the major components in the signal chain, the per channel power consumption of the ASCP for the medical grade application is approximately 830uW (rounded).

The number of channels (or the measurements *M*) required depends on the desired compression ratio (CR) and the reconstruction quality of the signal (PRD) (Refer to D.1.1 for the definition of PRD and CR). Previous research reported in [9] indicates that a PRD of less than 2% is considered '*very good*' quality of reconstruction and this corresponds to a CR of 20%. This corresponds to a M of 205 for a value of 256 for N, thus making the approximate estimate for the power of ACSP for this scenario to 170mW.

Scenario 2:

This scenario concerns about the ECG acquisition system for lifestyle applications where accuracy is not of paramount importance. Henceforth the specification is relaxed to the parameters in the below Table [10].

Parameter	Specification
Noise	3uV rms referred to input (RTI)
Bandwidth	128 Hz
Resolution	12 bits

Since the resolution for this application is moderate, one could use a SAR ADC which has the lowest FoM among the different architectures of ADCs. Typical values for FoM of a SAR ADC are in the range of 2fJ/conv-step to 30fJ/conv-step. For the conservative estimate we use 10fJ/conv-step for the analysis.

Similar reasoning can be extended to the choice of NEF for the amplifier given the relaxed specification. For the subsequent analysis we use a NEF of 2 (which is achieved for a current reuse architecture) to estimate power of FEA.

Also it has to be noted that the noise specification is still on a higher side for the given resolution and hence is modified to 1.5uV rms RTI. Employing the same procedure as followed in calculations in scenario 1, we get the ADC power as 10.4 nW per channel, FEA power as 0.26uW per channel and the mixer power as 1.5pW per channel. Hence the overall power consumption of the ACSP per channel for lifestyle application scenario is 0.3uW per channel.

For a lifestyle application a PRD value of 10% is deemed to be sufficient as it corresponds to a '*good*' reconstruction quality and the CR that can be achieved for this PRD level is around 70%. This corresponds to a value of 77 for M for a value of 256 for N. This yields an overall power consumption for the ACSP for this scenario as 23.1uW.

The summary of the calculations is listed in below Table

Application	Power consumption in uW per channel
Medical grade	830
Lifestyle	0.3

Appendix

In this section, we discuss the trade-offs involved between the analog and digital implementation of the class of systems that are realizable as filters. One of common functions in any signal processing chain is filtering and many of the signal processing algorithms can be recast in a way that they can be realized as filters. For example Fourier transform is usually realized in digital domain by using Fast Fourier Transform (FFT). It can also be realized in analog domain as well as a modulated filter banks. So, the question of implementation of a system in analog or digital domains is an interesting one, which we try to address here.

The usual problem with the analog circuits is the precision and is expressed using Signal to Noise Ratio (SNR). Attempts to increase the SNR results in higher power consumption (as indicated by NEF). A similar situation exists in digital implementation as well, where the power consumption scales with the number of bits and the number of bits (resolution) is a representative of precision. Hence one can intuitively think, that by examining the trade-offs between precision and power consumption in analog and digital domains, one could determine the optimal implementation. A similar study has been done by R. Sarpeshkar, the results of which are reproduced here for sake of completeness [12].
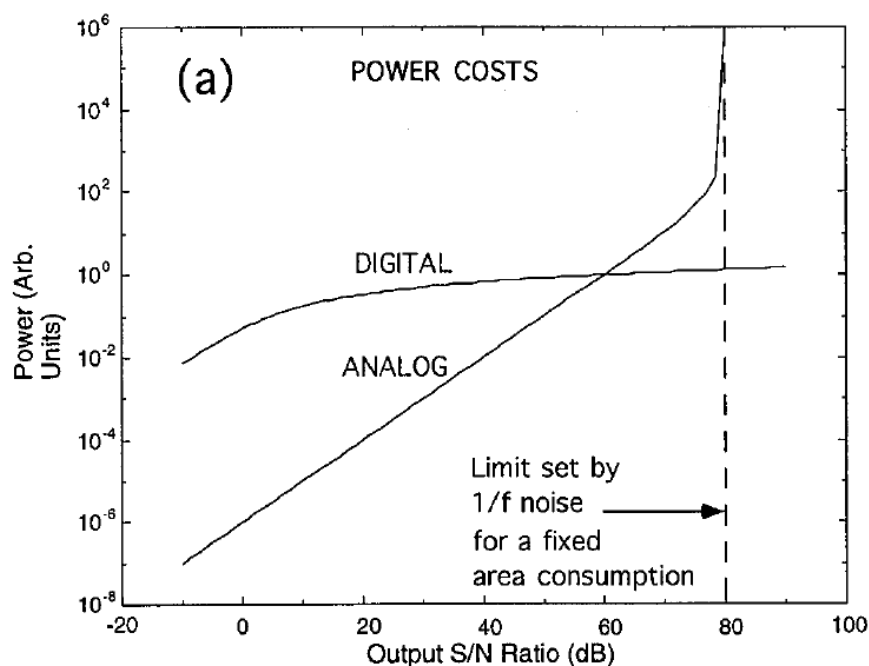
Figure 6: Power Vs Precision Tradeoff in Analog and Digital Systems

From the Figure 6, one can conclude that the implementing a system for precision less than 60 dB (~10 bits) in analog domain is efficient in terms of power and for higher precision a digital implementation is favourable. However, it must be mentioned that there are several subtle assumptions involved in the arriving at this conclusion, one of which is '*engineer ingenuity*'. In the computations we assume that the passive realization of a filter has the lowest power consumption for a given SNR and any active realization consumes much higher power. The ratio of the of the power consumed by the active implementation of a filter to the passive implementation is denoted by λ. A survey of

literature shows that the typical value of λ ranges from 100-1000 depending on the architecture and the creativity of the engineer (and hence the term *engineer ingenuity*) [11]. The results reported in [12] does not specify the value of the λ, but we re-formulate of the equations in terms of power, SNR and using a value of 100 for λ, we arrive at a conclusion that analog implementation is power efficient over digital implementations for a precision up to 68 dB (~11 bits). However, there does not seem to be a fundamental limit on how low the value of λ can be and until such a limit is established, the optimum power Vs SNR point has to be studied on case to case basis.

References

[1] Walt Kester, "Which ADC Architecture Is Right for Your Application?", Analog devices application note.

[2] B. Murmann, "ADC Performance Survey 1997-2013," [Online]. Available: http://www.stanford.edu/~murmann/adcsurvey.html.

[3] Pieter Harpe, Eugenio Cantatore, Arthur H. M. van Roermund, " A 2.2/2.7fJ/conversion-step 10/12b 40kS/s SAR ADC with Data-Driven Noise Reduction", ISSCC 2013, pp. 270-271.

[4] R. F. Yazicioglu, S. Kim, T. Torfs, H. Kim and C. V. Hoof, "A 30 uW Analog Signal Processor ASIC for Portable Biopotential Signal Monitoring," IEEE Journal of Solid-State Circuits, vol. 46, no. 1, pp. 209-223, Jan 2011.

[5] M. S. Steyaert and W. M. Sansen, "A micropower low-noise monolithic instrumentation amplifier for medical purposes", IEEE Journal of Solid-State Circuits, vol. 22, no. 6, pp. 1163–1168, 1987.

[6] Lei Liu, Xiaodan Zou, Wang Ling Goh and Minkyu Je, "Comparative Study and Analysis of Noise Reduction Techniques for Front-End Amplifiers", 13th International Symposium on Integrated Circuits (ISIC), 2011, pp. 555-558, 2011.

[7] R Muller, S Gambini and JM Rabaey, "A 0.013 mm^2 5μW DC-coupled neural signal acquisition IC with 0.5 V supply", IEEE Journal of Solid-State Circuits, vol. 47, no. 1, pp. 232-243.

[8] ANSI/AAMI-EC13, American National Standards for Cardiac Monitors, Hearth Rate Meters and Alarms, 2002.

[9] H. Mamaghanian, N. Khaled, D. Atienza and P. Vandergheynst. "Compressed Sensing for Real-Time Energy-Efficient ECG Compression on Wireless Body Sensor Nodes", IEEE Transactions on Biomedical Engineering, vol. 58, no. 9, pp. 2456-2466.

[10] Product sheet for nuubo nECG minder L1, [Online]. Available: http://www.nuubo.com/sites/default/themes/nuubo2/pdf/DATASHEETS_EN_minder.pdf

[11] Vittoz, E.A., "Low-power design: ways to approach the limits", Digest of Technical Papers. 41st ISSCC., 1994 IEEE International Solid-State Circuits Conference, 1994.

[12] R. Sarpeshkar, "Analog Versus Digital: Extrapolating from Electronics to Neurobiology," Neural Computation, Boston, 1998.