

# Generic vs custom; analogue vs digital: on the implementation of an online EEG signal processing algorithm

Alexander J. Casson and Esther Rodriguez-Villegas

**Abstract**—This paper quantifies the performance difference between custom and generic hardware algorithm implementations, illustrating the challenges that are involved in Body Area Network signal processing implementations. The potential use of analogue signal processing to improve the power performance is also demonstrated.

## I. INTRODUCTION

The intelligent sensor node for use in Body Area Networks is a well defined concept. A sensor is used to monitor a physiological parameter such as the ECG, EEG or temperature, and then local signal processing, on the sensor node itself, is used to effect immediate feedback and closed loop systems or to reduce the amount of data to store or transmit, allowing greater operational lifetimes. The heart of the sensor node thus becomes a signal processing algorithm.

It is well known that the power available from small, easily wearable, batteries is very limited and for a given size of battery historically the capacity has doubled only every 5–20 years [1]. The low power implementation of the signal processing algorithm thus becomes essential. For example, it has been shown that for wireless EEG applications compression or data reduction is required to make systems that operate from small batteries for a day or more feasible, and that this compression has a maximum power budget of only a few hundred micro-Watts [2].

To put this power budget in context, [3] investigates the hardware/software co-design implementation of a lossy EEG compression algorithm that compresses the data by 98%. However, the lowest power consumption found is 72 mW, over two orders of magnitude above the power budget for this level of compression for the compression to be power beneficial. It is clear that implementation of suitable algorithms will not be a trivial task.

This paper takes an example algorithm (based upon the one in [4]) and investigates how it can be implemented within the power budget available, which is found to be 96  $\mu$ W. The algorithm is for online data reduction in wireless EEGs for epilepsy diagnosis and is intended to be simple, giving a satisfactory performance level at the minimum power consumption, and so is a good candidate for investigating the power feasibility of such algorithms.

Two potential algorithm implementations are investigated here. Firstly the generic approach: essentially in the software domain using off-the-shelf processor solutions. Here the algorithm is in the digital domain. The second implementation

considered looks at the fully custom approach, but this time in the analogue domain.

It is found that it is not possible to achieve anywhere near satisfactory power performance using the generic approach. Furthermore, with modest assumptions, it is found that while a hardware digital domain implementation should be feasible, in the analogue domain it is feasible using just 24% of the available power budget. This result illustrates both the power challenges that must be faced for Body Area Networks to be truly realised, and that analogue signal processing, which at the fully custom level is not necessarily more specialised than a fully custom digital circuit, may be a preferable approach to tackling these power challenges.

A comparison such as the one carried out here will always contain a number of high level assumptions and it is obvious to some extent that dedicated circuits will always outperform more generic ones. Nevertheless, useful confirmation of this fact is presented. Also, the large discrepancy between the digital and analogue domains indicates that even relaxing some of the assumptions present an analogue approach is very competitive, and likely preferable.

The remainder of this paper is organised as follows: Section II summarises the EEG data reduction algorithm to be investigated and Section III derives the power budget available. The generic and digital domain power consumptions are then found in Section IV, the analogue estimates in Section V and the results discussed in Section VI.

## II. ALGORITHM OVERVIEW AND METHODS

The algorithm to be investigated here is a developed version of the one proposed in [4]. A high level overview of the procedure is given in Fig. 1. The overall aim of the algorithm is to reduce the system power consumption by recording only interictal (inter-seizure) epileptic events whilst not recording background signals. This gives a significant data reduction, reducing the storage or transmission power. As the algorithm aim is only data reduction, not event quantification, significant data reduction can still be achieved even with a number of false positives present [5].

The algorithm operates by processing all of the EEG channels recorded, a single channel at a time. A detection in any one channel causes all of the channels to record a section of EEG before and after the detection. Figures are presented here for a high quality 32 channel device (current portable EEG units typically have 16 channels) although this is essentially a variable. The  $z\beta$  parameters control the operation of the algorithm with  $\beta$  being user set and  $z$  being an automatically generated normalizing parameter to correct

The authors are with the Department of Electrical and Electronic Engineering, Imperial College London, SW7 2AZ {acasson, e.rodriguez}@imperial.ac.uk.

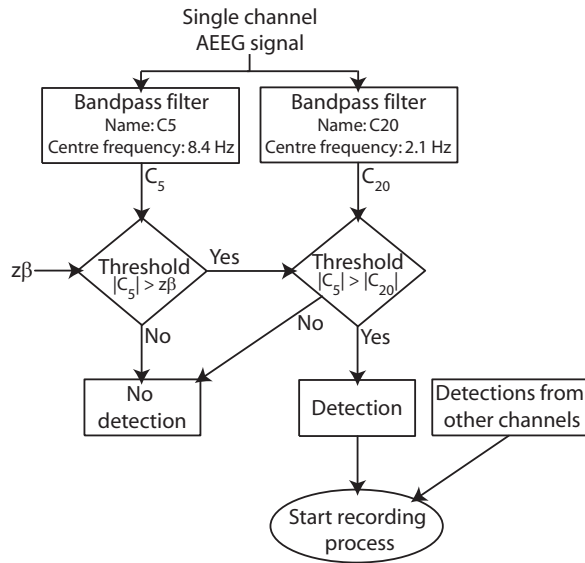


Fig. 1. A high level overview of the algorithm to be investigated. Detections in any monitored channel cause the algorithm to start recording.

TABLE I  
ANALOGUE AND DIGITAL ESTIMATES FOR A ONE CHANNEL ALGORITHM.

Component	Number required	Analogue power / W	Digital operations
Bandpass filters	2	200 n	46
Lowpass filter	1	20 n	7
Delay	2	200 n	4
Rectifier	1	140 n	1
Magnitude comparators	2	40 n	6
Multiplier	1	100 n	1
Switch	1	Negligible	1
Total	10	700 n	66

for broad level amplitude differences in different EEG traces. The effect of any data *buffering* during the recording process, so that data both before and after the detection point is recorded, is not considered here.

Although not all stages are shown explicitly in Fig. 1 a total of ten steps are performed each time the algorithm is run. These are listed in Table I. The overall intent is that this is a very simple signal processing algorithm that gives a satisfactory performance level at the minimum power consumption, and so is a good candidate for investigating the power feasibility of such algorithms.

In Section IV and Section V the power consumption of the algorithm is estimated by considering each of the ten steps in turn and, in the digital domain estimating the absolute minimum number of fundamental operations that are required, which is then linked to the power, and in the analogue domain by finding the power consumption of a typical, representative integrated component from our group's previous work all using the same process technology. Overall results are summarised in Table I.




Fig. 2. An example two channel wireless EEG system based upon [2]. The compression block can be in either the analogue or digital domain.

### III. POWER BUDGET

The power budget estimate here is based around the work from [2] and is explained by considering the two channel EEG acquisition system shown in Fig. 2, which can easily be extended to more channels. The basic architecture simply contains an instrumentation amplifier, an analogue to digital converter (ADC), a compression block (which contains the algorithm being investigated here) and a transmitter. The compression block can be freely placed either before or after the ADC for implementation in either the analogue or digital domain.

The power consumption of the entire system is given by

$$P_{\text{sys}} = NP_{\text{amp}} + NP_{\text{ADC}} + P_c + CP_t \quad (1)$$

where  $N$  is the number of channels,  $C$  is the compression ratio giving the ratio of the number of bits that are actually transmitted to the total number of bits if no compression was present,  $P_t$  is the power consumption of the transmitter, and the other three terms are the power consumptions of the amplifier, ADC and compression respectively. Just one compression stage is present to give the total power available for compression, but this can be broken down per channel if wanted.

If the transmitter has a power consumption of  $J$  Joules per bit  $P_t$  is given by

$$P_t = Jf_sRN \quad (2)$$

where  $f_s$  is the sampling frequency and  $R$  is the resolution in bits of the ADC.

If the system is operated with no compression stage present  $P_c = 0$  and  $C = 1$ . In order for the compression stage to be beneficial the following inequality must thus be satisfied:

$$P_c < JNf_sR(1 - C). \quad (3)$$

In practice, of course,  $P_c$  must be much lower than this to make a significant difference to the operating lifetime of the device, but this is not considered here. The results are intended to provide the upper bound on the power budget.

To minimize the power consumption  $f_s$  and  $R$  should be kept as low as possible. Typical values for the recording of clinical EEGs are given in [6] as  $f_s = 200$  Hz and  $R = 12$  bits. [2] gives a conservative estimate of  $J$ , which should be achievable in most situations as 50 nJ/bit, and a more speculative figure of 5 nJ/bit. This lower figure is used here so that any compression stage will not become obsolete if this

figure can be reliably realised. Finally,  $C$  is taken as 0.5 (a 50% data reduction), in line with the algorithm performance.

The power budget for the compression stage for its compression to be power beneficial is thus

$$P_c < 192 \mu\text{W} \quad (4)$$

although again, the actual power consumption may want to be an order of magnitude below this to be truly beneficial. Finally, as a high level assumption without explicit justification it is assumed that 50% of the power budget is reserved for the buffering of data before and after a detection is made (see Section II). This gives the end power budget for the 32 channel algorithm as  $96 \mu\text{W}$ .

#### IV. DIGITAL POWER ESTIMATE

##### A. Assumptions and fundamental limits

The analysis considered here is based around the use of essentially *off-the-shelf* micro-controller components. As a result the dynamic range of the signal processing is fixed regardless of that actually required. For example, if a 16 bit microprocessor is used calculations are assumed to be performed to 16 bit precision, even if this is not strictly necessary.

It is also noted that the power requirements of digital implementations are strongly dependent on the technology used to implement them. A high level model for the dynamic power (the power used while performing calculations) of a digital circuit assumes that power is only used to charge and discharge capacitive loads [1], [7]. The power consumption is then broadly given by

$$P = f C_T V_{DD}^2 \quad (5)$$

where  $f$  is the operating frequency,  $C_T$  is the total capacitance that is switched and  $V_{DD}$  is the supply voltage.

This basic model illustrates the high dependence of power in the digital domain on the supply voltage. Also, both  $C_T$  and  $V_{DD}$  tend to reduce as the technology feature size is reduced [7]. However, static power due to leakage currents tends to increase and it is possible for this to begin to dominate. This effect is not considered in the calculations below. It is noted, however, that for highly integrated systems the front-end, signal processing and transmitter will all be on the same chip. As a result it is not necessarily possible to arbitrarily scale the process technology to improve the performance of the signal processing as the extra leakage currents may limit the performance of the highly sensitive analogue front-end.

Finally, unless otherwise stated the calculations below are based around the number of instructions to be carried and so are independent of any duty cycling that may be present.

##### B. Operation count

The counts below illustrate the fundamental number of operations that are required to carry out each step in the algorithm each time it is run. These are intended to be very lower bound estimates and no weighting is applied for the relative complexity of different operations.

1) *Bandpass filters*: The generation of the two bandpass filter transfer functions required is detailed in the  $s$  domain in [8] where they have seven poles and two zeros. Converting these to the  $z$  domain using the MATLAB `c2d` function results in a filter with seven poles and six zeros. The filtering operation thus requires 12 multiplications and 11 additions or subtractions. These are all taken as elementary operations giving 23 operations per filter per filtering operation to be carried out.

2) *Lowpass filter*: Similar to the above, the  $z$  domain filter has two poles and one zero to implement. This thus requires 7 basic operations per filtering operation.

3) *Delay elements*: It is assumed that just two operations are required to implement a delay: one to store the current value and one to retrieve a previous value.

4) *Rectifier*: This is taken as just one operation to remove any sign bit which may be present.

5) *Magnitude comparators*: Similar to the above, two operations are assumed to remove any sign bits and then one further operation to perform the comparison. Thus three operations are needed in total.

6) *Multiplier*: This is taken as one basic operation.

7) *Switch*: Again this is taken as one operation.

##### C. Power consumption

The above estimates give a total of 66 operations to be performed each time the algorithm is run. Note that in practice it is highly likely that each algorithm block will require more than this bare minimum number of operations, each operation will correspond to a number of instructions and each instruction may take more than one clock cycle to execute. It may also be possible to perform more than one instruction per clock cycle. Preliminary results on a high performance Texas Instruments (TI) C6000 series Digital Signal Processing chip indicate that 2000 cycles per algorithm run are required, a factor of 30 more than the minimum number of operations, but this is not considered at this point.

An analogue version of the algorithm (see Section V) runs in continuous time and so to be comparable the digital one must be run each time a new sample is taken, at  $f_s$  which is taken as 200 Hz. It must also be run on all 32 channels meaning it must be run 6400 times a second giving 422 400 operations to be performed a second. This can now be related to the power consumption in several different ways. It is assumed here that each basic operation considered above is equivalent to one instruction to again give a lower bound solution.

Firstly, modern Intel processors, designed of course for computers rather than portable medical equipment, have an energy per instruction of approximately 10 nJ [7]. At 422 400 instructions per second this results in a power consumption of 4.2 mW, well above the  $96 \mu\text{W}$  power budget. To operate within this power budget only 200 pJ/instruction is available. To operate within a power budget of  $23 \mu\text{W}$  (see Section V) only 50 pJ/instruction is available.

Of course, these Intel processors aren't particularly suited to the portable situation in hand. As a more representative

TABLE II  
ACTIVE MODE POWER CONSUMPTION FOR A TI MSP 430  
MICROPROCESSOR WITH 2.2 V SUPPLY [11].

Clock speed / kHz	Current supply / $\mu\text{A}$	Power consumption / $\mu\text{W}$
1000	270	594
100	60	132
4	5	11

example the classic ARM7 processor has a performance of approximately 3.3 mW/MIPS (Millions of Instructions per second) on a 0.35  $\mu\text{m}$ , 3.2 V process [9]. In reality a performance of 230  $\mu\text{W}$ /MIPS is required, a factor of 14 improvement. The more modern ARM cortex-M3 processor is stated to be 70% (a 1.7 factor) more efficient [10], so this is still insufficient.<sup>1</sup>

As a final comparison, and one which allows a link to fully custom digital design, the popular for biomedical applications TI MSP430 microprocessor is considered, and this incorporates some of the overheads present in an actual implementation. Table II shows its typical power consumption [11]. Lower power operation is achieved by using lower clock speeds but it is clear that for operation at 96  $\mu\text{W}$  the device would have to be clocked at under 100 kHz and so there simply aren't enough clock cycles available to implement the required 422 400 operations per second.

It is thus clear that at this point in time it isn't realistic to use a generic approach to the signal processing. To link the current performance with that required the following argument is considered. The TI MSP430 is a 16 MIPS processor with a power consumption of 594  $\mu\text{W}$ , giving a performance of 37  $\mu\text{W}$ /MIPS. If the 430 duty cycled such that it has a 1 MHz clock and is in active mode 42% of the time to give sufficient clock cycles to perform the algorithm if the performance level of one operation per cycle can be achieved, the power consumption becomes 250  $\mu\text{W}$ . To fulfil the 96  $\mu\text{W}$  power budget an improvement by a factor of 2.6 to 14  $\mu\text{W}$ /MIPS is required. However, taking the factor of 30 from the preliminary DSP implementation this figure becomes approximately 0.5  $\mu\text{W}$ /MIPS which is inline with the performance of fully-custom ASICs which are of the order of 1  $\mu\text{W}$ /MIPS [12]. Thus, although the improvement factor with the assumptions present is fairly modest in practice it is highly likely that the performance of a fully custom implementation will be needed, but such an implementation should be feasible.

## V. ANALOGUE POWER ESTIMATE

### A. Assumptions and fundamental limits

The power consumption of an analogue implementation is taken by considering the power consumption of typical

<sup>1</sup>It is noted that [10] also gives the ARM7 power consumption as 0.28 mW/MHz and the cortex-M3 as 0.19 mW/MHz. If the processor could be operated at 422 kHz, with one high level operation per processor clock cycle these imply that acceptable power performance *may* be possible. Given the other analyses present here however this is currently deemed unrealistic.

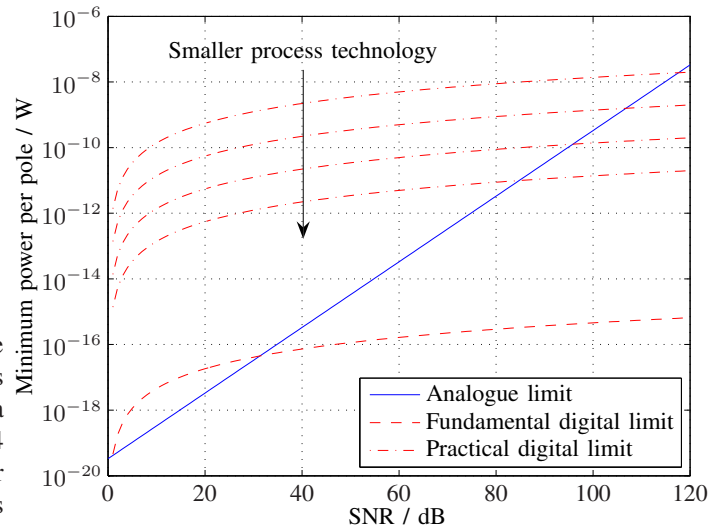


Fig. 3. The fundamental limits on the power requirements of analogue and digital signal processing versus dynamic range. From [15], [16].

components from our group and applying suitable safety factors. In general these components only have a dynamic range of around 45 dB (between 7 and 8 bits). This is considerably lower than the recommended EEG resolution of 12 bits [6]. However, it is noted that a typical diagnosis by a human from a digital EEG is performed with 16 channels on a screen with 1024 vertical pixels giving just 6 bits of resolution [13]. It is thus highly likely that this dynamic range is sufficient for the algorithm operation and this is confirmed by recent results by the authors [14].

This is significant due to the well known results from [15], [16] which are illustrated again in Fig. 3.<sup>2</sup> The results illustrate the fundamental limit for the power required for signal processing in the analogue and digital domains. It is derived principally for filter circuits (the core of the algorithm considered here) but is applicable in some other cases. It is found that the power consumption of a digital circuit is essentially independent of the dynamic range while that of an analogue circuit is a strong function of it. Thus at low dynamic ranges analogue circuits can give a better power performance. Of course, practical values are generally significantly above the fundamental limits, but this does give an expectation that dedicated analogue signal processing could outperform its digital counterpart. Also, it is noted from [15] that reducing the supply voltage of an analogue signal processing solution doesn't drastically improve the power performance as is the case for digital circuits.

### B. Block power estimates

1) *Bandpass filters*: The power estimate here is based upon the performance of the 6<sup>th</sup> order bandpass filter in [17] which has a power consumption of 70 nW. Taking an arbitrary safety factor of 1.4 to account for requiring a

<sup>2</sup>Note that absolute values in Fig. 3 depend upon the model used and the definition of dynamic range, but the qualitative result is unchanged.



higher order filter, needing lower centre frequency and not necessarily being able to replicate this performance, the power estimate is 100 nW per filter. Preliminary work on this filter indicates that this figure is achievable.

2) *Lowpass filter*: The low pass filter from [18] should be almost directly applicable and so no safety factor is taken giving an analogue power estimate of 20 nW.

3) *Delay elements*: In the analogue domain a delay element is essentially just a high order filter and so the estimate is again taken as 100 nW.

4) *Rectifier*: The power estimate is again taken from [18] which describes a roughly compatible rectifier with a 100 nW power consumption. Again taking a 1.4 safety factor the analogue power estimate is taken as 140 nW.

5) *Magnitude comparators*: In analogue it is possible to avoid the use of two rectifiers by using an inverting amplifier and two comparators. Based upon clocking down the comparator from [19], the complexity of this is assumed to be roughly equivalent to the low pass filter stage giving a power estimate of 20 nW.

6) *Multiplier*: An analogue multiplier is essentially the same as the transistor used in the filters required and so should be readily achievable within the 100 nW power budget given to them.

7) *Switch*: The power consumption of this simple stage is assumed to be negligible.

### C. Total power consumption

Combining all of the analogue figures gives a total power consumption estimate of 700 nW. The algorithm must be implemented for each channel and so for a 32 channel system the total power is 23  $\mu$ W; just 24% of the available power budget. This provides plenty of margin for error in the high level calculations that have been carried out here and also offers the possibility of longer system lifetimes because the compression occurs at such a low power level.

## VI. DISCUSSION AND CONCLUSIONS

This paper has investigated a signal processing algorithm with just ten steps, incorporating lower bound figures whenever possible and applying safety factors to ensure that the estimated power consumptions are achievable. Nevertheless, it is clear that a generic digital implementation is not currently feasible from a power perspective, and indeed the performance of a fully custom design is required before the power budget can begin to be realised. In contrast, however, the custom analogue solution requires just 24% of the power budget.

Whilst it is obvious to some extent that a specialised implementation will always outperform a generic one, the difference between the two performance levels has been illustrated, and is significant. Further demonstration of the potential use of analogue signal processing has also been given, quantifying its performance benefit.

Given this, in future work it is proposed to investigate the implementation of the algorithm in the analogue domain, against the traditional trend of digital solutions. This has

been shown to be capable of giving a very high performance level. While a digital implementation could be done at the Hardware Description Language level and re-synthesised to reap the benefits of technology scaling, this is not arbitrarily possible as doing so decreases the performance of the front-end system. Such a custom digital approach is also not intrinsically *easier* than the analogue approach.

Finally, it is likely that similarly limited power budgets apply to other Body Area Network applications and custom analogue approaches may be significant in meeting the power challenges presented.

## VII. ACKNOWLEDGEMENTS

The authors would like to acknowledge and thank Miss Elena Luna for her estimation of the clock cycles required to implement the algorithm on the TI DSP device.

## REFERENCES

- [1] A. Wang and A. Chandrakasan, "Energy-efficient DSPs for wireless sensor networks," *IEEE Signal Processing Mag.*, vol. 19, no. 4, pp. 68–78, 2002.
- [2] D. C. Yates and E. Rodriguez-Villegas, "A key power trade-off in wireless EEG headset design," in *IEEE EMBS NER*, Hawaii, May 2007.
- [3] A. Avila, R. Santoyo, and S. O. Martinez, "Hardware/software implementation of the EEG signal compression module for an ambulatory monitoring subsystem," in *IEEE ICCDCS*, Mexico, April 2006.
- [4] A. J. Casson, D. C. Yates, S. Patel, and E. Rodriguez-Villegas, "Algorithm for AEEG data selection leading to wireless and long term epilepsy monitoring," in *IEEE EMBC*, Lyon, August 2007.
- [5] A. J. Casson and E. Rodriguez-Villegas, "Data reduction techniques to facilitate wireless and long term AEEG epilepsy monitoring," in *IEEE EMBS NER*, Hawaii, May 2007.
- [6] G. Deuschl and A. Eisen, Eds., *Recommendations for the practice of clinical neurophysiology: Guidelines of the international federation of clinical physiology*, 2nd ed. Amsterdam: Elsevier, 1999.
- [7] E. Grochowski and M. Annaram, "Energy per instruction trends in intel microprocessors," *Technology@Intel Mag.*, March 2006.
- [8] A. J. Casson, D. C. Yates, S. Patel, and E. Rodriguez-Villegas, "An analogue bandpass filter realisation of the continuous wavelet transform," in *IEEE EMBC*, Lyon, August 2007.
- [9] S. B. Furber, *ARM system-on-chip architecture*, 2nd ed. Harlow: Addison-Wesley, 2000.
- [10] S. Sadasivan, "An introduction to the ARM Cortex-M3 processor," 2006. [Online]. Available: <http://www.arm.com/>.
- [11] Texas Instruments, "MSP430x22x2, MSP430x22x4 mixed signal microcontroller datasheet," 2007. [Online]. Available: <http://focus.ti.com/>.
- [12] F. Catthoor, "Ultra low power domain-specific instruction-set processor (ULP ASIP)," 2004, IMEC scientific report. [Online]. Available: <http://www.imec.be/>.
- [13] G. L. Krauss and R. S. Fisher, *The Johns Hopkins atlas of digital EEG: an interactive training guide*. Baltimore: Johns Hopkins University Press, 2006.
- [14] A. J. Casson and E. Rodriguez-Villegas, "Noise analysis and dynamic range considerations of an EEG spike selection algorithm," in preparation.
- [15] C. Entz and E. A. Vittoz, "CMOS low power circuit design," in *Desiging low-power digital systems: emerging technologies*, R. Calvin and W. Liu, Eds. Piscataway: IEEE press, 1996, pp. 80–133.
- [16] R. Sarpehshkar, "Analog versus digital: Extrapolating from electronics to neurobiology," *Neural Comp.*, vol. 10, no. 7, pp. 1601–1638, 1998.
- [17] P. Corbishley and E. Rodriguez-Villegas, "A nanopower bandpass filter for detection of an acoustic signal in a wearable breathing detector," *IEEE Trans. Biomed. Circuits Syst.*, vol. 1, no. 3, pp. 163–171, 2007.
- [18] P. G. Corbishley, "Ultra low power circuits for a miniature apnoea detection device," PhD thesis, Imperial College London, 2007.
- [19] E. Rodriguez-Villegas, "A 0.9 V offset compensated FGMOS comparator," in *IEEE ISCAS*, Kobe, May 2005.