

Design Of A 7-bit, 200MS/s, 2mW Pipelined ADC With Switched Open-Loop Amplifiers In A 65nm CMOS Technology

Carsten Wulff, Trond Ytterdal

Department of Electronics and Telecommunication, Norwegian University of Science and Technology, N-7491 Trondheim, Norway

Abstract—We present the design of a 7-bit 200MS/s pipelined ADC with switched open-loop amplifiers in a 65nm CMOS technology. As a result of turning off the open-loop amplifiers during sampling we reduce the power dissipation by 23%. The ADC achieves a SNDR of 40dB close to the Nyquist frequency, with a power dissipation of 2mW, which results in a Walden FOM of 0.13pJ/step and a Thermal FOM of 1.6fJ/step.

I. INTRODUCTION

Low resolution high speed ADCs have historically been Flash based architectures. The Flash ADC is a fast architecture due to its parallel nature. However, it is a very inefficient architecture. One of the reasons for its inefficiency is that it is not possible, with current processing technology, to get the input capacitance down to what is the minimum required by thermal noise. This is mostly due to parasitic capacitances from transistors and metal routing.

Before we begin our argumentation we should define what we mean with parasitic capacitances. We define parasitic capacitance as; any capacitance that is not required by the operation of the circuit.

One of the fundamental error sources that limit the performance of ADCs is the thermal noise power, which is usually represented as $\overline{V_{thermal}^2} = a \times kT/C$ where a is a constant, k is Boltzmann's constant, T is the temperature in Kelvin and C is the sampling capacitance.

Thermal noise power sets the lower limit for how much power we must dissipate to achieve a certain SNR. With respect to thermal noise power a certain SNR usually translates into a certain sampling capacitance.

If we assume a quantization noise power of $\frac{\overline{V_{LSB}^2}}{12} = V_{max}^2/(2^{2bits} \times 12)$, and we assume $\frac{1}{4}\overline{V_{LSB}^2} = \overline{V_{thermal}^2}$ and $a = 1$, the equation for the required sampling capacitor is

$$C = \frac{48kT2^{2bits}}{V_{max}^2} \quad (1)$$

If we use 6 bits, $k = 1.38 \times 10^{-23} J/K$, $T = 353K$ and $V_{max} = 0.4V$ the required sampling capacitance is 6fF. This, of course, assumes that other concerns like mismatch or unwanted effects from parasitic capacitances does not come in to play.

Parasitics at a node can reach 10fF in current nanoscale technologies, as a consequence the parasitics can be larger than required sampling capacitor at the 6-bit level. This is one

reason why a Flash ADC loses when it comes to efficiency. A 6-bit Flash ADC without averaging or interpolation has 64 comparators connected to the input, each of which has possibly 10fF input capacitance. Thus a Flash ADC can have 600fF input capacitance, which is two orders of magnitude larger than the required sampling capacitance.

If we look at the figure of merit (FOM) of 6-bit ADCs, and higher resolution ADCs, using the Walden FOM [1] given by

$$FOM = \frac{P_{diss}}{2^{bits} f_s} \quad (2)$$

where P_{diss} is the power dissipation and f_s is the sampling frequency. And the Thermal FOM given by

$$FOM = \frac{P_{diss}}{2^{2bits} f_s} \quad (3)$$

We get Figure 1, with the gray diamonds representing Walden FOM and the black triangles representing Thermal FOM. The data for Figure 1 was collected from ADCs published in the Journal of Solid State Circuits in the years 1975-2005. According to the Walden FOM there are 6-bit ADCs that are just as good as the 15-bit ADCs, since they have the same figure of merit. However, the Walden FOM is an empirically deduced FOM, and it has come under some scrutiny in the recent years. A more theoretically correct FOM is the Thermal FOM.¹

The argument for why the Thermal FOM is more correct is as follows. Assume a thermal noise limited ADC, where the power dissipation is proportional to the sampling capacitance. If we increase the resolution by one bit, we can see from (1) that the sampling capacitance quadruples, and through this the power dissipation quadruples. However, the Walden FOM only allows a doubling of the power dissipation for each bit of resolution added. This leads to an unfair FOM for high number of bits and a lenient FOM for low number of bits.

Two alternatives to Flash ADCs have received some attention in recent years. The first is successive-approximation (SAR) ADCs and the second is pipelined ADCs. Both have in recent papers achieved good results.

In [2] they presented a 6-bit 600MS/s time interleaved asynchronous successive-approximation ADC, with a Walden FOM of 0.22pJ/step and a Thermal FOM of 5.7fJ/step. In

¹As far as we know the FOM has not yet been given a name, so the name Thermal FOM is not an official name. It is, however, a descriptive name.




Fig. 1. Walden FOM versus Thermal FOM as a function of bits for ADCs published in the Journal of Solid State Circuits 1975-2005. Thermal FOM in black and Walden FOM in gray.

[3] they presented a 6-bit pipelined ADC with open-loop amplifiers achieving a Thermal FOM of 59fJ/step. Note that the best 6-bit ADC in Figure 1, is [4] with a Thermal FOM of 16.38fJ/step, which is three orders of magnitude worse than the best > 14-bit ADCs. And the best 7-bit ADC is a 100kS/s SAR [5] with a Thermal FOM of 1.89fJ/step, which is two orders of magnitude worse than the best > 14 bit ADCs.

Our goal for this design was to optimize the FOM for a low resolution ADC at a reasonable speed. We choose to design a pipelined ADC and placed it at the conservative speed of 200MHz. That is, the speed is conservative if we compare to the speed to [3] or [2]. Usually, low resolution ADCs are used in the GHz range, so we assume that the this pipelined ADC would be used in a time interleaved architecture. Knowing that we would be unable to use a 6fF sampling capacitor we opted for 7-bit resolution, since the thermal noise power would anyway be low because of the larger sampling capacitor, and adding one more stage does not increase the power significantly. To keep parasitic capacitances from transistors as low as possible we chose to use a 65nm CMOS technology. The pipelined ADC architecture is explained in the following section with results of simulations presented in Section III.

II. PIPELINED ARCHITECTURE

The ADC was designed in a 65nm low power CMOS technology with low threshold voltage (lvt) transistors. The architecture of the differential pipelined ADC is shown in Figure 2. The ADC has five 1.5-bit pipelined stages and a three level flash at the end. Each stage has a three level analog to digital converter (SADC) and a multiplying DAC (MDAC). The MDAC has a gain of two.

One of the alternatives to the traditional operational amplifiers in MDACs is an open-loop amplifier, like a common

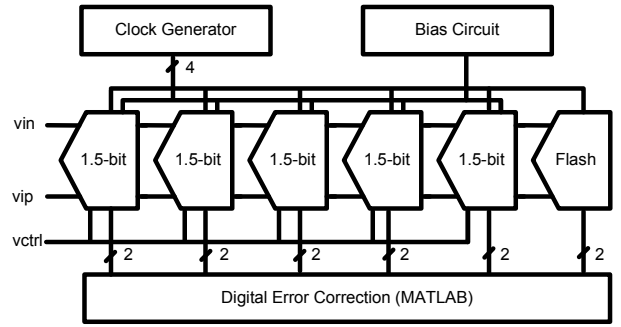


Fig. 2. Architecture overview of the 7-bit Pipelined ADC with open-loop amplifiers

source amplifier, with a gain of two. This technique has been used with success in a 12-bit pipelined ADC [6] and 6-bit pipelined ADC [3]. Our design is based on [3]. The MDAC architecture can be seen in Figure 3. Figure 3 shows stage one and stage two. Stage one is in the amplifying phase, the SADC has made its decision, and the control signals t_0 , t_1 and t_2 control transmission gates that connect one of the two sampling capacitors to high reference, common mode or low reference. The control signals t_0 , t_1 , t_2 and the transmission gates implement the DAC. Stage two is shown in the sampling phase. Here both capacitors are connected to the input through transmission gates controlled by clock ip_1 . Each stage needs three clock phases, p_1 , p_{1a} and p_2 , where p_{1a} is slightly advanced over p_1 , and is used to sample the input when it is quiet. p_1 and p_2 are non-overlapping clock phases. Stage two uses ip_1 , ip_{1a} and ip_2 , where $ip_1=p_2$ and $ip_2=p_1$. All in all we need four clock phases for the complete pipelined ADC, p_1 , p_{1a} , p_2 and p_{2a} . The open-loop amplifier is marked by x_2 in Figure 3.

A. Open-Loop Amplifier

A detailed schematic of the open-loop amplifier can be seen in Figure 4. It is a differential common source amplifier with resistive load. The resistors are each $4k\Omega$, and we assume that they will be calibrated at startup, so the resistor value is constant over process. In this differential amplifier it is important that the common mode variation is low. If not, the common mode voltage will drift as the signal propagates through the pipelined stages, and may even turn later stages off. To achieve low common mode variation we use a replica bias that keeps the total current, $I_{tot} = I_{R1} + I_{R2}$, equal to two times I_{bias} over process, voltage and temperature variations. Through this the common mode voltage is determined by $V_{cm} = V_{dd} - I_{tot}/2 \times (R1 + R2)/2$ which keeps it constant over process and temperature variation.

The transistors $M5/M6/M11/M12$ are twice the size of $M7/M8/M11/M12$, as a consequence $I_{tot} = 2I_{bias}$. We choose the common mode to be 0.6V, to get larger overdrive on the input transistors, and the swing to be $\pm 0.2V$.

The gain of the common source amplifier, if we disregard the source degeneration, is given by $A_o = g_{m1}R_1$ The gain




Fig. 3. Stage 1 and stage 2 in the pipelined ADC

will vary over process and temperature because of changes in g_m , we compensate for some of this change by varying the v_{ctrl} voltage of the source degeneration transistor, which in effect changes the effective source degeneration resistor and in turn changes the gain of the amplifier. If we include the source degeneration the gain expression becomes

$$A_o = \frac{g_{m1}R_1}{1 + \frac{g_{m1}}{2g_{ds13}}} \quad (4)$$

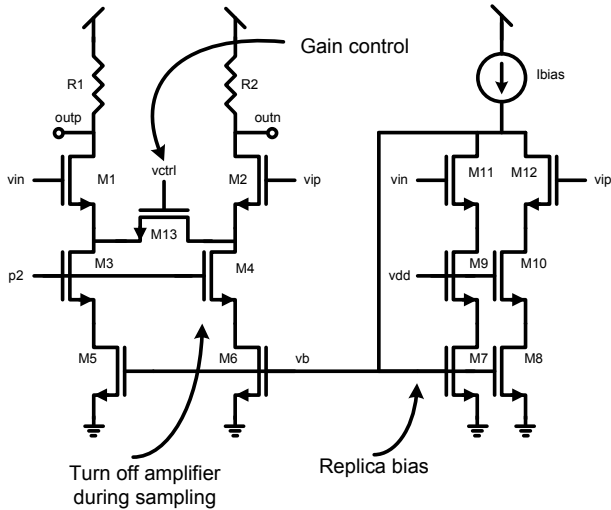


Fig. 4. Open-Loop Amplifier

In [3] they used a replica MDAC stage configured in a feedback loop to control the v_{ctrl} voltage such that the MDAC stage has a gain of two. Since we only do simulation we have

not included the replica stage, the v_{ctrl} voltage is changed manually for each corner simulation.

The gain of the MDAC is also affected by the input capacitance of the amplifier as described in [3].

Our main contribution to reduce the power of the pipelined ADC is the transistors M3/M4, these turn off the the amplifier during sampling phase when it is not needed. Because the drain and source nodes of M3/M4 are low impedance, and the bias voltage, v_b , stays constant, the amplifier turns on quickly.

In our design the bias current is $I_{bias} = 100\mu A$, which makes the total current consumption of a MDAC stage $300\mu A$, ergo the power dissipation of five MDAC stages should be $1.5mW$, assuming we use a 1V supply. This is if we leave the amplifiers on all the time. By turning off the amplifiers during sampling the power dissipation is reduced to $P_{mdacs} = 5 \times 100\mu A \times 1V + \frac{1}{2} \times 5 \times 200\mu A \times 1V = 1mW$, so we would expect an improvement of 0.5mW when the amplifiers are turned off during sampling.

With this low current in the amplifiers, the input capacitance of the next stage must be low. The sampling capacitors (C1, C2) are chosen at 50fF. The reason for choosing such a large value compared to the required, is not capacitor matching concerns, which do not come in to play at 7-bit, but rather the parasitic capacitance from the amplifier input, which reduces the gain of the MDAC. We compensate for some of this reduction in gain with the v_{ctrl} voltage.

Two other circuits dissipate power in the pipelined ADC, the SADCs and the clock generator/buffers. The size of the clock buffers are mainly determined by the load of the SADCs and routing capacitance. The power dissipation in the SADCs are determined by matching concerns of the comparators. The comparators used are the so called Lewis-Gray dynamic comparators introduced in [7].

B. Clock Generation

Most switched capacitor circuits use two non-overlapping clocks to control the charge transfer, but since we use bottom plate sampling there is an advanced phase 1 that transitions just before phase 1. The advanced clock phase reduce the problem of input dependent charge injection from the input switches.

We use NMOS inputs in the comparators due to the high common mode voltage, thus the comparators use a clock that samples on the rising edge. To avoid distributing inverted clock phases we invert the clocks of the transmission gates, which means that the transmission gates are on when the clock is low, and off when the clock signal is high. The clock buffers were scaled to drive 6 SADCs, transmission gates, amplifier turn-off control signal and a 100fF capacitance was added at the output of each clock buffer to model the the parasitic capacitance added by metal routing.

III. RESULTS OF SIMULATION

The ADC was simulated at transistor level using Eldo from Mentor Graphics. A common mode of 0.6V, and a swing of 0.2V, as previously mentioned, resulted in a differential peak to peak of 0.8V with a 1V supply. Both the references and

inputs are assumed to be buffered off-chip, the buffers are not included in the simulation. Results of Eldo simulations were post-processed in MATLAB, where an FFT was performed and signal-to-noise-and-distortion (SNDR) was extracted. In all simulations an input signal frequency close to the Nyquist frequency was used, and an input signal amplitude of 0.8 times full scale range.

Mismatch simulations were performed in the typical corner at a temperature of 27 degrees Celcius. A 2^7 point FFT was used to estimate the SNDR and 101 simulations were run to get an estimate of the standard deviation of the SDNR due to mismatch. The mean SNDR was 40dB and the standard deviation was 1.2dB.

Four process corners (fast, slow, fast-slow, slow-fast) and three temperature corners (0° , 27° , 80°) were simulated. When we vary vctrl to compensate for the threshold voltage changes, the standard deviation in SDNR over process corners and temperature corners is 2.35dB, with the worst corner being slow process and low temperature. With constant vctrl the standard deviation increases to 3.45dB.

Figure 5 shows an 2^{10} point FFT of the output from a transient simulation with noise in the typical corner.




Fig. 5. A 1024 point FFT of the ADC output from a transient noise simulation. The harmonics of the fundamental are marked with diamonds.

The total power dissipation for the simulated ADC in typical corner (typical process, 27° and 1V supply) is 2mW. If we leave the amplifiers turned on during sampling it dissipates 2.6mW, with the increase being close to the expected 0.5mW. Table I summarizes the achieved performance in the typical corner.

The Walden FOM of the ADC is 0.13pJ/step and the Thermal FOM of 1.6fJ/step. The improvement is almost a factor four compared to the Thermal FOM of 5.7fJ/step from [2], bearing in mind that they have proven silicon, and that the difference might be eaten up by layout parasitics. In addition, we operate at a lower speed and a higher SNDR, which makes it more straightforward to achieve a good figure of merit. However, [2] was a SAR while our ADC is a pipelined ADC, making the point the two architectures can compete on equal grounds with respect to figure of merit at reasonable speeds.

TABLE I
PERFORMANCE SUMMARY OF THE 7-BIT PIPELINED ADC

Technology	65nm LP CMOS
Input Voltage Peak-to-Peak	0.8V
Supply Voltage	1V
Sampling Frequency	200MHz
SNDR	40dB
ENOB	6.3dB
Power Dissipation	2mW
Walden FOM	0.13pJ/step
Thermal FOM	1.6fJ/step

IV. FUTURE WORK

Low resolution pipelined ADCs with open-loop amplifiers make for an interesting architecture, it may well be the most efficient, straightforward way to implement high speed low resolution ADCs at the present time. More work is needed to try to reduce the sampling capacitance, and we believe that cutting the parasitic capacitances down, through architecture or technology changes, is the most effective way to increase effectiveness of high speed low resolution ADCs.

V. CONCLUSION

We presented the design of a 7-bit 200MS/s pipelined ADC with switched open-loop amplifiers in a 65nm CMOS technology. As a result of turning off the open-loop amplifiers during sampling the power dissipation was reduced by 23%. The ADC achieved a SNDR of 40dB close to the Nyquist frequency, with a power dissipation of 2mW, which resulted in a Walden FOM of 0.13pJ/step and a Thermal FOM of 1.6fJ/step.

ACKNOWLEDGMENTS

Financial support from the Norwegian Research Council through the project Smart Microsystems for Diagnostic Imaging in Medicine (project number 159559/130) and the project ASICs for Microsystems (project number 133952/420) is gratefully acknowledged.

REFERENCES

- [1] R. H. Walden, "Analog to digital converter survey and analysis," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, pp. 539–550, 1999.
- [2] S.-W. M. Chen and R. W. Brodersen, "A 6-bit 600ms/s 5.3mw asynchronous adc in 0.13- μ m cmos," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2669–2679, 2006.
- [3] D.-L. Shen and T.-C. Lee, "A 6-bit 800ms/s pipelined a/d converter with open-loop amplifiers," *IEEE J. Solid-State Circuits*, vol. 42, no. 2, pp. 258–268, 2007.
- [4] J. H.-C. Lin and B. Haroun, "An embedded 0.8 v/480 uw 6b/22 mhz flash adc in 0.13-um digital cmos process using a nonlinear double interpolation technique," *IEEE J. Solid-State Circuits*, vol. 37, no. 12, 2002.
- [5] M. D. Scott, B. E. Boser, and K. S. J. Pister, "An ultralow-energy adc for smart dust," *IEEE J. Solid-State Circuits*, vol. 38, no. 7, 2003.
- [6] B. Murmann and B. E. Boser, "A 12-bit 75-ms/s pipelined adc using open-loop residue amplification," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, 2003.
- [7] T. B. Cho and P. R. Gray, "A 10 b, 20 msample/s, 35 mw pipeline a/d converter," *IEEE J. Solid-State Circuits*, vol. 30, no. 3, 1995.