A high bandwidth power scaleable sub-sampling 10-bit pipelined ADC with embedded sample and hold

Imran Ahmed
Department of Electrical and Computer Engineering
University of Toronto
Toronto, Canada
imran@eecg.toronto.edu

David A. Johns
Department of Electrical and Computer Engineering
University of Toronto
Toronto, Canada
johns@eecg.toronto.edu

Abstract—**A pipelined ADC architecture for use in sub-sampled systems which has its power scaleable with down sampled bandwidth is presented. Using a technique developed to eliminate the front end sample and hold, a power savings of >20% is achieved compared to a previous design. A technique to improve the settling behavior of Rapid Power on Opamps is also presented. Measured results in 1.8V 0.18µm CMOS verify the removal of the front end sample and hold does not cause gross MSB errors for input frequencies higher than 267MHz. With f_s=50MS/s, for f_{in}=79MHz the SNDR is 51.5dB, and with f_s=4.55MS/s for f_{in}=267MHz the SNDR is 52.2dB.**

I. Introduction

With the growing demand for mobile systems which can provide multi-bandwidth and multi-standard services in a single solution, research in reconfigurable systems has recently garnered much attention. Many mobile systems use a sub-sampled ADC for down conversion to baseband or low IF as shown for one path of an IQ solution in Fig. 1 so as to eliminate a mixer at the input. By directly feeding the modulated signal to the ADC however the input bandwidth of the ADC is significantly increased.

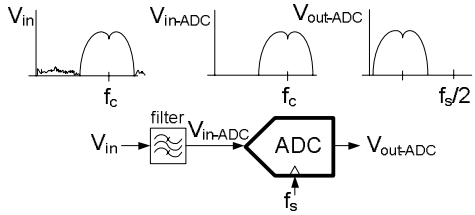

Fig. 1: Down conversion using a sub-sampled ADC

As shown in Fig. 2, in a pipelined ADC, the first pipeline stage is sensitive to skew between the MDAC and sub-ADC, where if enough skew is present gross MSB errors occur resulting in a dramatic reduction in ADC accuracy. Skew is caused by mismatch in time constants between input signal paths and increases with input frequency, hence is a critical issue for sub-sampling systems. A front end Sample and Hold (S/H) is commonly used before the first pipelined stage to eliminate skew. The front end S/H however must have a kT/C noise floor lower than that of the pipelined ADC following it, hence the S/H is typically is a significant power consumer in pipelined ADCs. Previous publications have eliminated the front-end S/H by relying on the redundancy of the pipelined stage e.g. [1], [2]. In a 1.5b/stage architecture the comparator offset in the sub-ADC can be as large as $V_{ref}/4$ (where V_{ref} is the maximum peak voltage of the input). Thus so long as the difference in input operated on by the MDAC and sub-ADC of Fig. 2 is less than $V_{ref}/4$, the effect of skew appears as an offset on the sub-ADC comparator and the front end sample and hold can be eliminated without any further modification to the ADC. Assuming a sinusoidal input to the ADC with the maximum amplitude V_{ref}, input frequency f_{in}, and assuming no offset in the sub-ADC comparators it can be determined the maximum allowable skew in a 1.5b stage is $(8\pi f_{in})^{-1}$ [2]. For f_{in}=270MHz, the maximum skew allowable is only 140ps. When the offset errors of the comparators are included, the allowable skew is much smaller. Clearly approaches which rely on the redundancy of the input stage require extremely carefully matched layout in sub-sampled systems where the input frequency is very large.

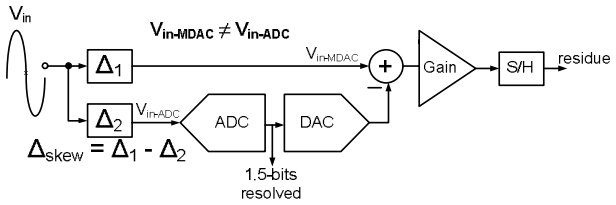

Fig. 2: skew in signal path in pipeline stage

In this work a power scaleable ADC optimized for sub-sampled systems with large input bandwidths is proposed. By using a modified architecture for the first pipeline stage which does not require a front-end S/H nor carefully matched signal paths to prevent gross errors in the first stage, the power of the ADC is reduced by 20% compared to that in [3]. This work also presents a technique to improve the settling behavior of Rapid Power on Opamps. Measured results from a prototype fabricated in a 1.8V 0.18µm CMOS process show the ADC of this work to have an SNDR of 51.5dB at f_{in}=79MHz for f_s=50MS/s, and SNDR of 52.2dB at f_{in}=267MHz for f_s=4.55MS/s.

II. Design Architecture

A. Elimination of front-end Sample and Hold

A conventional 1.5b stage used in a pipelined ADC is illustrated in Fig. 3. During Φ_1 the input is sampled on capacitors C_1 and C_2, and during Φ_2 a gain of two is implemented by discharging the charge stored in C_1 to C_2, and DAC operation by connecting V_{DAC} to a voltage set by the sub-ADC.

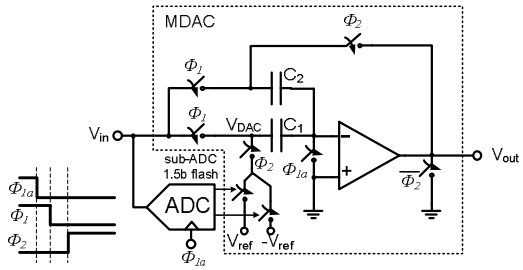

Fig. 3: Conventional 1.5b pipeline stage

In this work the first 1.5b pipeline stage is modified and an additional clock phase introduced as shown in Fig. 4. Like a conventional pipeline stage, during Φ_1 the input is sampled on C_1 and C_2 however in the modified pipeline stage when Φ_2 switches from low to high and Φ_{2D} is low, V_{DAC} is set to a high impedance and C_2 is connected between the output and input of the opamp. Due to charge conservation a held value of the sampled input is thus produced at the opamp output as shown in Fig. 5. Thus with the modified approach during the time labeled t_{delay} in Fig. 4 the output of the first stage can be connected to a 1.5b flash ADC. When Φ_{2D} subsequently goes high the 1.5b flash ADC resolves its input and sets V_{DAC} to the appropriate DAC voltage. Thus the modified pipeline stage implements the same functionality as a conventional pipeline stage, however the approach is not sensitive to skew at the input as by using the embedded sample and hold of a pipeline stage the sub-ADC operates on the same input that is sampled by the MDAC regardless of input frequency.

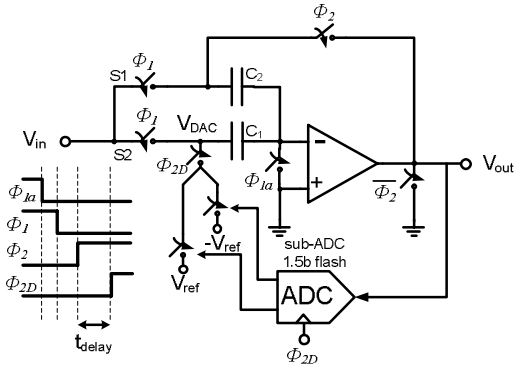

Fig. 4: Modified pipeline stage to eliminate front-end S/H

Since the flash ADC only requires an input that is >1.5b accurate, the output of the opamp during t_{delay} is not required to fully settle to generate the correct outputs from the flash ADC hence t_{delay} can be a small fraction of the total available settling time. Although the power of the opamp is slightly increased due to a reduced settling time to implement a gain of 2, the overall power of the ADC is significantly reduced as the power hungry front-end S/H is eliminated. The technique can also be applied to multi-bit pipeline stages, noting that prior approaches which relied on redundancy to eliminate the front end S/H have an even smaller allowable skew hence demand an even more meticulously matched layout than discussed in section I. It should be noted that the technique to remove the front end S/H although independently derived for this work is similar to that published in [4]. However the results of this work show that this approach can be applied to much higher input bandwidth as well as a 66% increase in sampling rate.

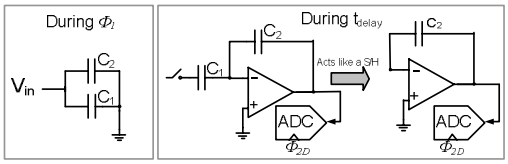

Fig. 5: Operation of modified pipeline stage

B. *Power Scaleability*

The power of the sub-sampling ADC in this work is scaled by using the Current Modulated Power Scaling (CMPS) technique described in [3]. CMPS achieves power scaling by digitizing the analog input within a fixed time interval t_{ON}, then powering off the ADC for t_{OFF} until the next input sample is required to be sampled. Thus with CMPS the settling times for each pipeline stage are constant for different sampling rates where settling time is set by the clock f_{sm} in Fig. 6 [3].

From [3], when using CMPS the ADC effectively operates as an algorithmic ADC, limiting the maximum sampling rate to $1/t_{lat}$=4.55MS/s in this work, where t_{lat} is the latency of the ADC. Current scaling is used for power scalability for sampling rates not covered with CMPS [3] (i.e. between 4.55-50MS/s). By using CMPS with current scaling a very wide power scalable range can be realized, which is highly desirable in sub-sampled systems, as the system could be configured for a low bandwidth signal such as voice, or configured for a high bandwidth signal such as video (which has a bandwidth orders of magnitude larger than voice), while having the power scale with bandwidth.

III. Circuit Implementation

A. *Pipeline Architecture*

The architecture of the pipelined ADC of this work is shown in Fig. 6. The first stage is as shown in Fig. 4, and all remaining stages are standard 1.5b stages. Stages 3-9 are identical to those used in [3]. Stage 2 is a standard 1.5b stage but has the same sized sampling capacitors and opamp as stage 1.

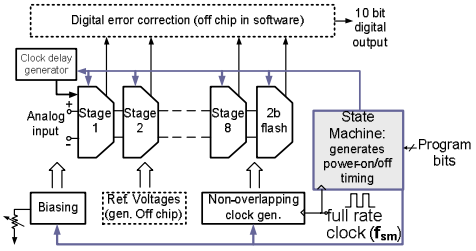

Fig. 6: ADC architecture

With the front end S/H removed in this work, the total sampling capacitances of the first two pipelined stages were reduced by 40% compared to [3] (from 940fF to 580fF in stage 1, and 2). In this work the opamp of stage 1 was conservatively designed such that t_{delay} was 20% of the total settling time. However t_{delay} was made programmable so that the lowest t_{delay} possible without gross MSB errors could be measured. To meet the same settling requirements as [3], the

opamps of stage 1 and 2 were made 33% smaller than those of [3]. Thus although less settling time is available in the first pipeline stage, overall power is still reduced without a front end S/H since the thermal noise floor of the pipeline stages can be made higher to maintain a fixed input referred noise floor. To enable a large input bandwidth, bootstrapped switches [5] were used as the input switches S1, S2 in Fig. 4.

B. Rapid Power on Opamp

In [3] Rapid Power on Opamps were used in each pipeline stage to realize a power scaleable architecture, and are shown in Fig. 7. Like conventional gain-boosted opamps, the Rapid Power on Opamp requires the unity gain frequency of the gain booster opamps (opamps A_p, A_n in Fig. 7) to be higher than the 3dB frequency of the main opamp but lower than the unity gain frequency of the main opamp [6]. Furthermore the loop formed by the gain boosters (labeled L_{Ap} in Fig. 7) is also required to be stable. Rapid Power on Opamps power on and completely off each clock cycle, hence the inputs to the gain booster opamps A_p, A_n effectively see a step function every clock cycle. Thus unlike conventional gain-booster circuits where the phase margin of the gain booster opamps need not be very high (as the inputs typically do not see large voltage jumps), Rapid Power on Opamps require the loop L_{Ap} to have a very good phase margin. In [3] standard folded cascode PMOS-input opamps as shown in Fig. 8 were used for the gain booster opamps A_p. The phase margin of the loop labeled L_{Ap} in Fig. 8 is limited by the second pole of the loop which occurs at the PMOS current mirror node labeled V_{p2} in Fig. 8. PMOS transistors were sized ~5x larger than the NMOS transistors to maintain a similar overdrive voltage as the NMOS transistors for a fixed current density. As such the parasitic gate-source and gate-drain capacitance on node V_{p2} is large and sets the second pole of the loop to a low frequency degrading phase margin of the loop. To improve the phase margin of loop L_{Ap}, the p-input gain booster opamp was modified by implementing the current mirror in the opamp with NMOS transistors instead of PMOS transistors [7] as illustrated in Fig. 8. By performing the mirror operation with NMOS transistors, the capacitance on the mirror node is reduced by ~2.5x and the second pole pushed to a much higher frequency significantly improving phase margin. Simulation results show that the phase margin of loop L_{Ap} is improved from 56º to 72º by only changing the location of the mirror node. By switching the location of the current mirror node from PMOS to NMOS transistors, the slew rate of the gain boosting opamp is reduced by 2x. Although this increases the power-on time, the significant benefit of a more stable power-on transient makes it a favorable trade-off.

C. Generation of Delayed Clock Φ_{2D}

The delayed clock edge Φ_{2D} of Fig. 4 was generated using a chain of four current starved inverters as shown in Fig. 9. Current starved inverters were chosen to allow t_{delay} to be widely tuned for different sampling rates. Also by varying the off chip reference current I_{ref}, the dependence of ADC SNDR vs. t_{delay}, could be measured.

The current starved inverters were designed such that only one clock edge was delayed. The current source transistors were connected between the drains of the PMOS and NMOS inverter transistors. This was done so that when the input to e.g. the first inverter switches from low to high, the current source transistor MB1 switches from operating in cut-off to the active region forcing the inverter output to discharge with rate set by the bias current of MB1. Hence the discharge rate at the output of the inverter is a strong function of the biasing current of MB1 allowing wide variation in the delay of the inverter. If however the current sources were connected to the source nodes of the inverters (as is often done and shown in Fig. 10), when the input switches from low to high in e.g. the first inverter, the current source initially starts off in triode, and discharges the pre-charged output until the current source is biased in the active region as shown in Fig. 10. I.e. the discharge rate of the output also becomes a function of the rise time of the input, reducing the control I_{ref} has on the delay of the current starved inverter, hence reducing range of delay values possible by varying I_{ref}.

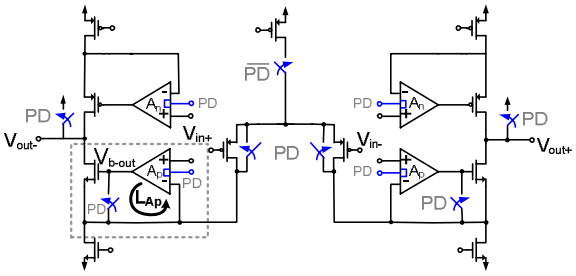

Fig. 7: Rapid Power-on Opamp

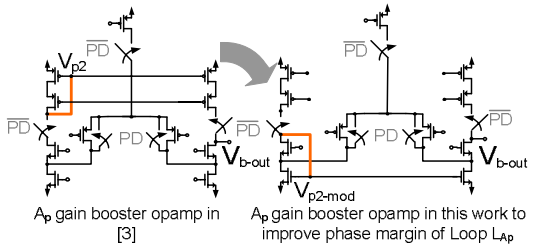

Fig. 8: Modified gain booster, A_p, for improved PM

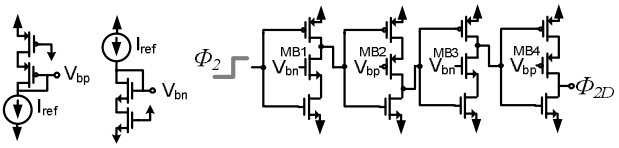

Fig. 9: Current starved inverter chain to generate Φ_{2D}

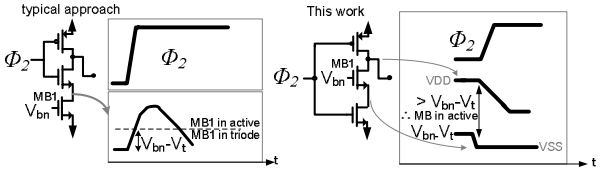

Fig. 10: Impact of where current source is connected

IV. MEASURED RESULTS

A prototype of the ADC of this work was fabricated in a 1.8V, 0.18µm CMOS process (shown in Fig. 11); the core area was 1.1 mm x 1.1 mm. The maximum input swing was 1.6V

p-p. Fig. 12 shows the SNDR of the ADC versus input frequency for different power scaled sampling rates. The ADC remains fully functional for input frequencies larger than 267MHz – frequencies which prior techniques using redundancy would require very well matched layout (<<140ps skew). For f_s=50MS/s and f_{in}=79MHz the SNDR is 51.5dB. For lower sampling rates the input bandwidth is increased, e.g. for f_s=4.55MS/s, the SNDR is 52.2dB for f_{in}=267.1MHz. Fig. 13 shows the FFT of the ADC output for f_s=50MS/s and 4.55MS/s. We note in this work the S/H removal technique is demonstrated for input frequencies ~4x larger than [4]. It is noted that the measurements for f_s=4.55MS/s are for the case when CMPS is enabled such that the settling time in each opamp as well as t_{delay} is the same as the case when f_s=50MS/s [3]. Also note that a single sinusoidal input was used but distortion products are captured in the results due to aliasing from sub-sampling.

Fig. 14 shows the power of the ADC versus sampling rate. The power of the ADC at f_s=50MS/s was 27mW, >20% lower than the 35mW of [3]. It is noted in this work the 27mW includes an additional bias circuit (1mW) to improve the robustness of the system, as well as clock delay generator (0.5mW). Lower sampling rates with correspondingly lower power can be realized by increasing t_{ON} while current scaling the ADC [3].

Fig. 15 shows the SNDR of the ADC vs. % t_{delay} is of the total available sampling time. MSB errors only occur when t_{delay} is <10% of the available settling time. Thus the technique to remove the front-end sample and hold described in this work does not require the first stage opamp to be significantly increased in power to maintain settling accuracy. Furthermore the fact that the SNDR only degrades for t_{delay} larger than 30% of the settling time indicates the power of the first stage opamp could easily be further reduced.

V. Conclusions

A power scaleable ADC for sub-sampled systems with a large input bandwidth was described. Using a technique to remove the front end sample and hold, a power savings of >20% was realized. A method to improve the settling behavior of Rapid Power on Opamps was also presented. Measured results from a 1.8V 0.18µm CMOS prototype show the ADC to achieve more than 51dB SNDR for input frequencies larger than 79MHz for f_s=50MS/s and 267MHz for f_s=4.55MS/s.

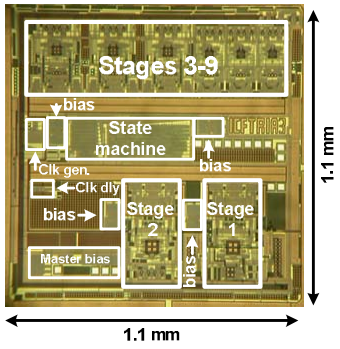

Fig. 11: Micrograph of 1.8V, 0.18µm CMOS chip

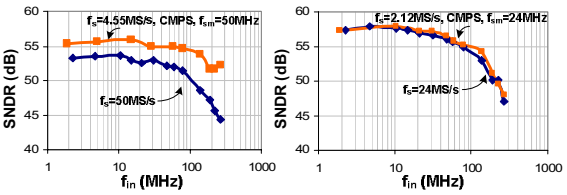

Fig. 12: SNDR vs. f_{in} for different f_s

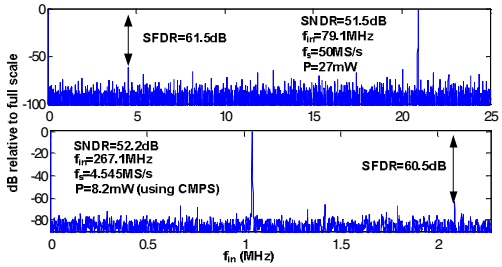

Fig. 13: FFT of ADC output for f_s=50, 4.55MS/s

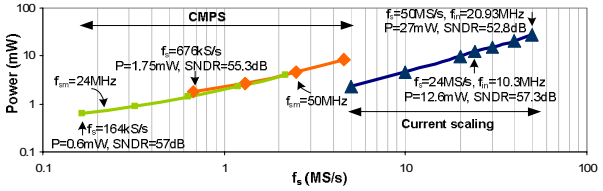

Fig. 14: Power vs. sampling rate

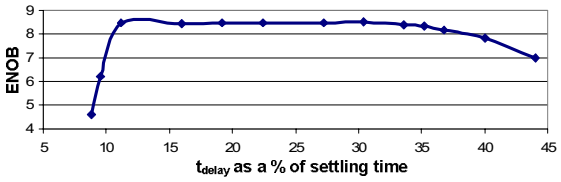

Fig. 15: ENOB vs. t_{delay} as % of settling time (f_s=50MS/s)

References

[1] K. Gulati et al, "A Highly Integrated CMOS Analog Baseband Transceiver With 180 MSPS 13-bit Pipelined CMOS ADC and Dual 12-bit DACs," *IEEE Journal of Solid State Circuits*, vol. 41, pp. 1856–1866, Aug. 2006

[2] A A. Ali et al, "A 14-bit 125 MS/s IF/RF Sampling Pipelined ADC With 100 dB SFDR and 50 fs Jitter", *IEEE Journal of Solid State Circuits*, vol 41, pp. 1846-1855, Aug. 2006

[3] I. Ahmed, D. Johns, "A 50MS/s (35mW) to 1kS/s (15µW) power scaleable 10b Pipelined ADC Using Rapid Power On opamps and Minimal Bias Current Variation," *IEEE Journal of Solid State Circuits*, vol 40, pp. 2446-2455, Dec. 2005

[4] J. Li, X. Zeng, L. Xie, J. Chen, J. Zhang, Y. Guo, "A 1.8-V 22mW 10-bit 30-MS/s Subsampling Pipelined CMOS ADC," *IEEE Custom Integrated Circuits Conference*, Sept. 2006, pp. 513-516

[5] A.M. Abo, P.R. Gray, "A 1.5 V, 10-bit, 14MS/s CMOS Pipeline analog-to-digital converter" *IEEE, VLSI Circuits*, June 1998, pp. 166-169

[6] K Bult, G.J.G.M. Geelen, "A Fast settling CMOS Op Amp for SC Circuits with 90-dB DB Gain", *IEEE Journal of Solid-State Circuits*, vol 25, pp.1379-1384, December 1990

[7] R.J. Baker. CMOS Circuit Design, Layout, and Simulation. IEEE Press/ Wiley Interscience, New Jersey, 2005