

A 1.35 GS/s, 10 b, 175 mW Time-Interleaved ADC Converter in 0.13 μm CMOS

Simon M. Louwsma, A. J. M. van Tuijl, *Member, IEEE*, Maarten Vertregt, *Member, IEEE*, and Bram Nauta, *Fellow, IEEE*

Abstract—A time-interleaved ADC is presented with 16 channels, each consisting of a track-and-hold (T&H) and two successive approximation (SA) ADCs in a pipeline configuration to combine a high sample rate with good power efficiency. The single-sided overrange architecture achieves a 25% higher power efficiency of the SA-ADC compared with the conventional overrange architecture, and look-ahead logic is used to minimize logic delay in the SA-ADC. For the T&H, three techniques are presented enabling a high bandwidth and linearity and good timing alignment. Single channel performance of the ADC is 6.9 ENOB at an input frequency of 4 GHz. Multichannel performance is 7.7 ENOB at 1.35 GS/s with an ERBW of 1 GHz. The FoM of the complete ADC including T&H is 0.6 pJ per conversion step. An improved version is presented as well and achieves an SNDR of 8.6 ENOB for low sample rates, and, with increased supply voltage, it reaches a sample rate of 1.8 GS/s with 7.9 ENOB at low input frequencies and an ERBW of 1 GHz. At $f_{in} = 3.6$ GHz, the SNDR is still 6.5 ENOB, and total timing error including jitter is 0.4 ps rms.

Index Terms—Analog-to-digital converter (ADC), bandwidth mismatch, high-speed sampling, jitter, pipeline, SAR, successive approximation ADC (SA-ADC), time-interleaved, time-interleaving, timing, timing alignment, track-and-hold (T&H).

I. INTRODUCTION

ATREND in receiver design for digital TV, satellite receivers, and set-top boxes is the move towards software-defined radios, where the embedded analog-to-digital converter (ADC) is moved closer to the antenna. Such ADCs require 8–10 bits of resolution, a large bandwidth to enable subsampling/down-conversion, and limited power consumption of a few hundred milliwatts to be able to embed the ADC with the digital baseband processing in a single IC. The time-interleaved ADC architecture shown in Fig. 1 offers the combination of good power efficiency with high speed [2], [5], [6]. Key aspects are high-speed sampling, excellent matching between channels, and power efficiency of the complete system. This paper extends the results presented in [5] and [9] with in-depth analysis and the presentation of an improved design.

This paper is organized as follows. First, the time-interleaved track-and-hold (T&H) is discussed in Section II, with channel matching issues between the T&H channels in Section III. Second, the sub-ADC architecture is presented in Section IV,




Fig. 1. Time-interleaved ADC architecture.

followed by calibration aspects and measurement results in sections V and VI. The paper ends with the presentation of an improved design and its performance in Section VII and conclusions in Section VIII.

II. TIME-INTERLEAVED T&H

A time-interleaved T&H needs a signal bandwidth per channel far beyond the sample rate of an individual channel. Moreover, it requires matching between channels. The T&H presented here is a successor to the one presented in [3] and has improved bandwidth, linearity, and channel matching. Innovations presented are: an improved buffer design, a bandwidth enhancement technique, and a low skew clock-driver, and it includes gain and offset calibration.

A. Number of Channels

By time-interleaving multiple ADCs, the operation speed is increased by the interleaving factor [7]. A tradeoff exists between this factor and the input capacitance. When the interleaving factor is too high, this capacitance is also too high, and a power-hungry buffer is required to drive the capacitance with sufficient bandwidth and linearity [1]. We use a moderate interleaving factor of 16 so the load can be driven by a 50 Ω source and no buffer is required [3]. For our specifications, this interleaving factor requires sub-ADCs with a sample rate of about 100 MS/s and about 50 dB SNDR. Such ADCs can be implemented with a good power efficiency.

The timing diagram of the time-interleaved T&H is shown in Fig. 2. At each falling edge of the master clock (CLK), one of the T&H channels goes from track mode to hold mode and takes a sample of the input signal. Conventionally, the track time is equal to the hold time. In this design, we use only one clock cycle of the master clock for tracking. This is sufficient and has two advantages: 1) more time is available for the ADCs to perform their conversion and 2) the input capacitance is reduced because now only one instead of eight sample capacitors is connected to the input simultaneously.

Manuscript received August 24, 2007; revised November 4, 2007.

S. M. Louwsma, A. J. M. van Tuijl, and B. Nauta are with the IC Design Group, CTIT Research Institute, University of Twente, 7500 AE Enschede, The Netherlands (e-mail: S.M.Louwsma@alumnus.utwente.nl).

M. Vertregt is with NXP Semiconductors, High Tech Campus WY2, 5656 AE Eindhoven, The Netherlands.

Digital Object Identifier 10.1109/JSSC.2008.917427




Fig. 2. Timing diagram of the time-interleaved T&H.





Fig. 3. Overview of 1 of the 16 T&H channels; S2 is open during tracking to increase the buffer bandwidth.

Fig. 4. High-linearity input buffer thanks to constant V_{DS} of the PMOST.

The basic schematic of one T&H channel is shown in Fig. 3. The circuit is (quasi-)differential, but only a single-ended version is shown. Bootstrapping of the sample switch [11] is used to reach good linearity. The purpose of switch S2 will be discussed in Section C.

B. Linear Buffer

The schematic of the unity-gain T&H buffer of Fig. 3 is shown in Fig. 4. It is in fact a P-type source-follower, with an additional N-type source-follower aiming to keep the drain-source voltage of the pMOS transistor constant. In modern submicron processes, the output resistance of minimum length MOSFETs is very small and nonlinear, resulting in reduced gain and distortion. The second source-follower decreases the variation in V_{DS} , which increases the effective output resistance of the pMOS transistor and results in improved gain and linearity.

Furthermore, compared with a conventional source-follower, the input capacitance is not increased like in a cascoded source-follower [6], but it is decreased: The gate-drain (overlap) capacitance is effectively lowered since the drain terminal has roughly the same phase and amplitude as the gate terminal. A low and linear input capacitance is important to avoid distortion of high-frequency input signals.

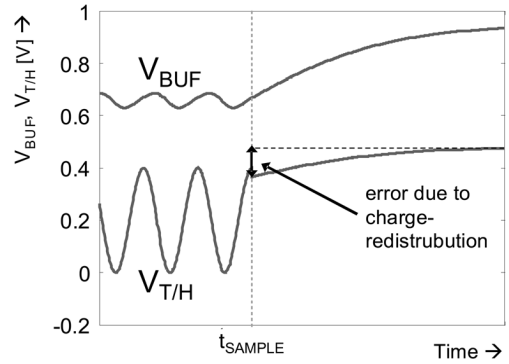


Fig. 5. Sampling a high-speed input signal with limited buffer bandwidth.

C. High-Bandwidth Sampling Technique

For a stand-alone T&H and a buffer with first-order settling behavior, the bandwidth requirement for the buffer for settling is: $BW > (n+1) \cdot \ln(2) \cdot 2 \cdot f_S / 2\pi$ with n bit resolution and half of the sample period for settling. For $n = 10$, $BW > 4.9f_{Nyquist}$. An input buffer with this bandwidth even tracks input signals at the Nyquist frequency closely. For time-interleaved T&Hs, the bandwidth requirement for settling is relaxed by the interleaving factor. The bandwidth requirement for a time-interleaved T&H is: $BW_{INT} > (n+1) \cdot \ln(2) \cdot 2 \cdot f_S / 2\pi \cdot ifac$ where $ifac$ is the interleaving factor and again with half of the sample period for settling. With $n = 10$ and an interleaving factor of 16, the bandwidth requirement is: $BW > 0.3f_{Nyquist}$. When a buffer with minimal bandwidth for settling is used to save power, the buffer output no longer tracks input signals at the Nyquist frequency and a large attenuation and phase shift is present. Now, the problem as shown in Fig. 5 arises: during tracking, the buffer output V_{BUF} cannot follow the input signal $V_{T\&H}$, and, at the sample moment (t_{SAMPLE}), output signal V_{BUF} is not yet fully settled. After the sample moment, the buffer output V_{BUF} will slowly settle to its final value. During this settling, charge redistribution between: 1) the nonlinear parasitic capacitance between the input and output of the buffer and 2) the sample capacitor causes distortion in the voltage of the sample capacitor $V_{T\&H}$ and the buffer output V_{BUF} .

To avoid distortion, the buffer bandwidth can be increased, but this will increase the power consumption significantly. Instead, we introduce a switch S2 between the buffer output and the ADC as shown in Fig. 3 [9]. In track mode, this switch is open and the load capacitance of the buffer is small. Hence, the buffer bandwidth is high and output V_{BUF} can now follow the input $V_{T\&H}$, as shown in Fig. 6. In this case, no distortion due to charge redistribution occurs.

When the ADC is connected at $t = t_{SWITCH}$, the buffer output will first make a step to the previous sample value still present on the ADC input capacitance. Then, the buffer will charge the ADC load to the new sample value (see Fig. 6). Charge redistribution after $t = t_{SWITCH}$ also causes a signal-dependent step in $V_{T\&H}$, marked by S. This seems to cause distortion; however, as V_{BUF} settles to its final value, the process of charge redistribution is reversed and $V_{T\&H}$ returns to its initial undistorted value.

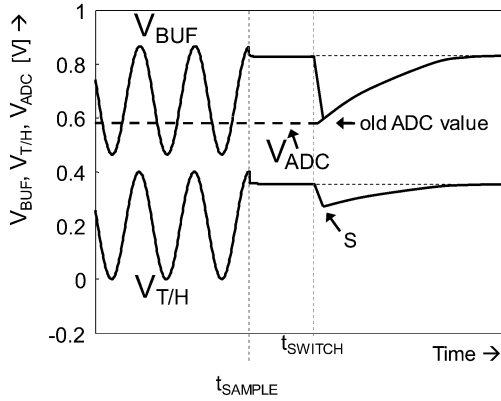


Fig. 6. Sampling a high-speed input signal with enhanced buffer bandwidth.

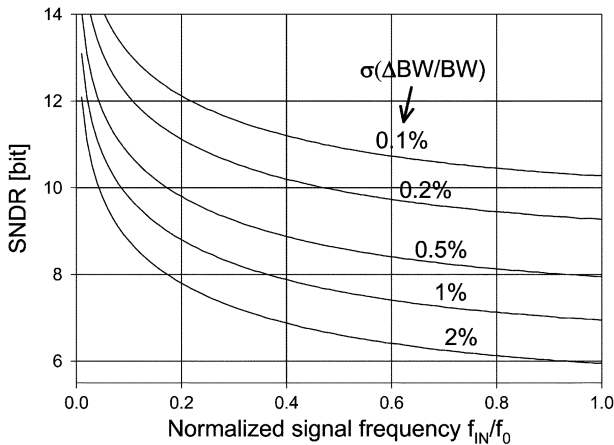


Fig. 7. SNDR as a function of the normalized input frequency for different values of $\sigma(\Delta BW/BW)$. f_0 is the nominal channel bandwidth.

III. CHANNEL MATCHING

In a time-interleaved T&H, differences between channels should be sufficiently small: offsets between channels cause tones at multiples of f_S/N , while differences in gain, bandwidth, or timing result in tones at multiples of $f_S/N \pm f_{in}$, where N is the interleaving factor [7], [10].

Reducing channel offsets by straightforward device sizing to reach the required accuracy conflicts with speed and power constraints [3], therefore channel offsets are made adjustable in this design. Channel gain is adjustable as well to correct for errors stemming from mismatch in the T&H buffers and the sample process. The implementation is discussed in Section V.

A. Bandwidth Matching

Bandwidth mismatch between channels causes frequency-dependent differences in gain and phase [10]. In Fig. 7, the SNDR due to bandwidth mismatch as a function of the normalized input frequency is shown for different values of $\sigma(\Delta BW/BW)$. When the channel bandwidth is equal to the Nyquist frequency ($f_S/2$), σ (gain) should not exceed 0.1% for an SNDR of 10 bits and input frequencies up to $f_S/2$.

To check whether this is feasible, we need to take a few assumptions. A sample capacitor of 150 fF is sufficient to limit kT/C noise below a 10 bit level. When a bandwidth of 1 GHz

is required, the switch resistance R_{ON} should be 1 k Ω . A minimum-length N-type MOS-switch with a width of 1 μm is just sufficient in the 0.13 μm process. The standard deviation of the resistance mismatch $\sigma(\Delta R_{ON}/R_{ON})$ is 3.5%, while capacitor matching is much better. $\sigma(\Delta BW/BW)$ is therefore 3.5% as well, which is much larger than the required 0.1%. Improving resistance matching by device scaling is impractical, because $\sigma(\Delta R_{ON}/R_{ON})$ only scales with $1/\sqrt{W}$. However, when the width of the switch is increased, the bandwidth is increased as well and the impact of the mismatch at the signal frequency becomes lower. When, e.g., $\sigma(\Delta BW/BW)$ is 1% and the bandwidth is chosen to be ten times larger than the highest signal frequency, an SNDR of almost 10 bits is possible (when only considering bandwidth mismatch). For this reason, the channel bandwidth has to be chosen larger than what is required when just looking at signal attenuation.

B. Timing Matching

Calibration of timing mismatch requires high-frequency test signals and complicated calibration algorithms. By careful design, we aim for a timing alignment within the required accuracy. For a large number of channels, timing offsets between channels can be approximated as a Gaussian distribution. In this case, the SNR caused by timing offsets is:

$$\text{SNR}_{\Delta t} = 1/\{\sigma(\Delta t) \cdot 2\pi \cdot f_{in}\}.$$

For an SNR of 50 dB and an f_{in} of 1 GHz, the required timing offset between channels should be smaller than 0.5 ps rms. In [2], a technique to prevent timing errors in a time-interleaved T&H is presented. It uses a front-end sampling switch which is closed only half of the period of the master clock. A disadvantage of this method is the decrease in bandwidth, which makes it unsuitable for high signal frequencies and high interleaving factors. In this design, we achieve good timing alignment by using a master clock [4] to synchronize the different sampling instants and matched lines to distribute clock and input signals to the channels (see Fig. 14).

In applications where supply noise may degrade performance, current-mode logic (CML) is commonly used because it generates little supply noise. CML uses differential signaling, with a typical signal swing of half of the supply voltage. To convert the CML master clock into a full-swing signal suitable for the sample switch, a conversion circuit is needed. In a previous chip, we used a conventional solution consisting of a differential pair with a current mirror on top, followed by a transimpedance stage and a buffer stage. Using this circuit, we measured a timing misalignment of 6 ps rms, which is much too high for the target specification. Therefore, a new circuit topology is proposed which minimizes the path from the common master clock to the sample switches. The circuit is shown in Fig. 8 and operates as follows: the T&H is put in track mode by the bootstrap circuit and, at the end of the track mode ($t = t_2$), node V_S is left floating by the bootstrap circuit, and for further bootstrapping we rely on parasitic capacitances. Now, to switch into hold mode, node V_S has to be discharged to ground rapidly: transistors P1 and N1 take care of this. Assume switch S1 is closed and $V_{CP} < V_{CN}$, so node V1 is at ground potential and P1 is nonconducting. In this state,




Fig. 8. Proposed CML-to-SE conversion circuit, together with waveforms.

switch S1 is opened, without influencing the potential of V1. When the differential voltage of the master clock ($V_{CP} - V_{CN}$) becomes larger than the threshold voltage V_T of P1, P1 will start to conduct and node V1 will be charged to V_{CP} , as shown in Fig. 8. This will make N1 conducting, discharging node V_S rapidly and putting the circuit into hold mode. This is the only time-critical event in the T&H.

Advantages of this solution are: 1) only the mismatch of P1 and N1 influences the skew, so the complete “spread budget” can be spent in these transistors; 2) both differential clocks are used and the effective slope is doubled, which halves the influence of the threshold voltage variations of P1; and 3) by minimizing the number of transistors between the clock input of the chip and the sample switch, jitter is minimized.

The expected timing misalignment is 0.45 ps rms. This value is derived from multiplying the (simulated) switching slopes by the $\sigma(\Delta V_T)$ of the respective transistors.

In a time-interleaved T&H, the channels should sample one after the other, with a delay of one clock-period (see Fig. 2). Each period, only one of the channels should switch into hold mode. This is accomplished by applying clock gating to the circuit, and this does not influence the performance. Note that the load for clocks CP and CN is not symmetrical and causes some imbalance. However, this does not degrade the timing alignment.

IV. SUB-ADC ARCHITECTURE

The 16 sub-ADCs in this design require a sample rate of about 100 MS/s and an SNDR of 50–55 dB. To sufficiently reduce quantization noise, we choose a resolution of 10 bits. For these specifications, various architectures exist, such as pipeline and two-step converters. The successive approximation (SA) ADC (SA-ADC) architecture is chosen here for its high power efficiency; it uses only one comparator, and high-gain amplifiers—a necessity in pipeline and two-step architectures—are not required, which makes it suitable for nanometer-scale technologies. However, its operation speed is limited at high resolution: for an n bit converter, n iterations are required. In order to exploit the advantages of an SA-ADC in this design, three techniques are used to get a high sample rate and good power efficiency: 1) two SA-ADCs are combined in a pipeline configuration; 2) single-sided overrange technique; and 3) look-ahead logic.

A. Pipelining

An overview of the sub-ADC architecture which is used 16 times in the complete ADC is shown in Fig. 9. It consists




Fig. 9. Overview of the sub-ADC architecture (1/16 of the total ADC).

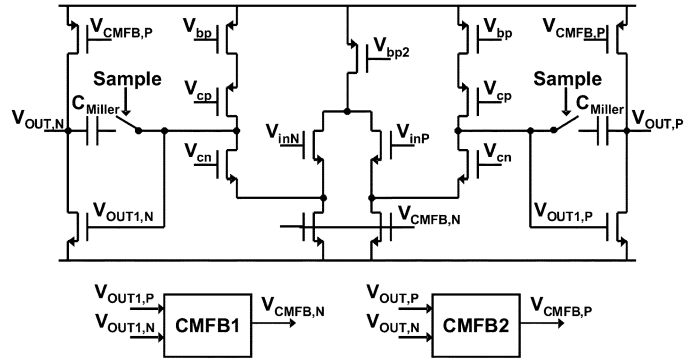


Fig. 10. Opamp schematic with Miller capacitance only connected during sample mode.

of an interleaved T&H section, a first 6 bit SA-ADC, a digital-to-analog converter (DAC), an inter-stage amplifier, and a second 6 bit SA-ADC. To simplify debugging, the outputs of both SA-ADCs are not combined, but they are directly made available off-chip. All signaling is pseudo-differential. The amplifier has an implicit T&H, such that both SA-ADCs have a full period to do their conversion. Compared with a single 10 bit SA-ADC, their requirements are relaxed: less accuracy is required and fewer steps are needed, resulting in more time per step.

The gain of the amplifier is 16, such that the bits of the second ADC have to be shifted $^2 \log(16) = 4$ positions to the right in order to get the same weights for both ADCs. The overall resolution is 10 bit and the overrange is 4 LSBs of the first ADC. This large overrange relaxes the requirements on the inter-stage amplifier significantly, because after amplification by 16 the residue signal is nominally only a quarter of the range.

In pipeline converters with a high resolution, the number of bits in the first stage is increased to lower the requirements on the MDAC. In such converters, mostly a flash ADC is used, but its resolution is limited due to its high input capacitance. In our design, we use an SA-ADC, which only has a low input capacitance and allows us to use 6 bits in the first stage, largely reducing the MDAC (amplifier) requirements.

The amplifier is realized by a two-stage opamp in switched-capacitor configuration and cancels opamp offset. To stabilize the common-mode signal, switched-capacitor common-mode feedback is applied (see Fig. 10). It has two modes: 1) sample mode, where the feedback factor is 1, requiring a large Miller capacitor for stability and 2) amplify mode, where the feedback factor is 1/16 and no frequency compensation is needed, but a




Fig. 11. Circuit for controlling the current and common-mode voltage of the resistor ladder.

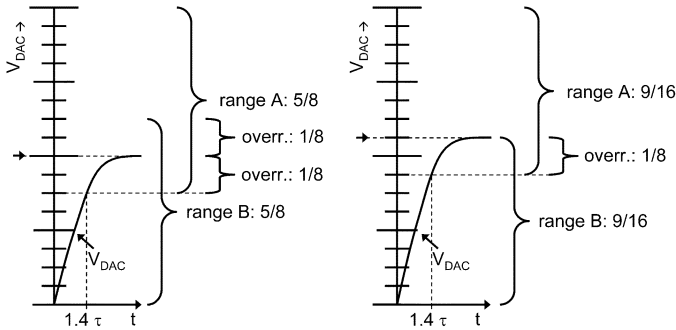


Fig. 12. Overrange techniques. (a) Conventional. (b) Single-sided.

16 times higher gain–bandwidth product is required. This conflict between bandwidth and stability is resolved by connecting the capacitor only during sample mode.

The DAC is a resistor ladder with switches and the ladder is shared between all channels, avoiding differences between channels. The current and common-mode level of the ladder are controlled by the circuit shown in Fig. 11. The reference voltage and DC current are controlled externally, however no external voltage buffers are required.

B. Single-Sided Overage Technique

An SA-ADC operates in a loop: the comparator makes a decision, the control logic determines the next DAC level, the DAC settles to the next value, and so on. All three actions have to be completed within one clock cycle. For an n bit converter, n iterations are required. A technique to reduce the delay caused by the DAC settling is discussed next, followed by a technique to reduce the delay of the logic.

For a conventional SA-ADC, the DAC settling error should be smaller than $1/2$ LSB. For an RC -limited DAC, the time required for settling is: $t_{\text{SETTLE}} > n \cdot \tau \cdot \ln(2)$, where n is the number of bits and τ is the settling time constant of the DAC. For a 6 bit converter, 4.2τ of DAC settling is required in the first clock cycle. The settling time can be reduced by employing the principle of overranging [8]. When an overrange of $1/8$ of the range is used as shown in Fig. 12(a), the DAC settling can be reduced to only 1.4τ . The next range has the size of $5/8$ of the previous range (the radix is $8/5$) and $^2 \log(8/5) = 0.68$ bit is resolved each cycle, requiring 9 cycles to reach 6 bits of accuracy.

The overrange technique assumes overrange on both sides of the comparison level. An RC -limited DAC, however, does not show overshoot during transitions and the DAC error is only

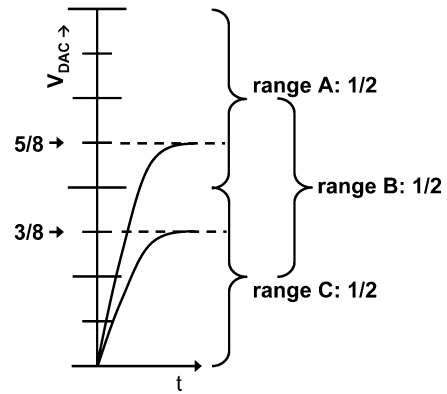


Fig. 13. Architecture with two comparators and two DACs.

TABLE I
OVERVIEW OF SA-ADC ARCHITECTURES

Architecture	# steps	settling time	total time	energy per conversion
conventional	6	4.2τ	25.2τ	25.2
1/8 overrange	9	1.4τ	12.6τ	12.6
single-sided overrange	7	1.4τ	9.8τ	9.8
2 comparators	6	1.4τ	8.4τ	16.8

due to incomplete settling. Note that this is also true when all ladders are combined, as long as the impedance of the ladder is much smaller than the impedance of the DAC switches. An overrange is therefore only needed at the side from which the DAC is settling. This is schematically shown in Fig. 12(b) for a rising DAC signal (for a falling DAC signal, the diagram is vertically flipped). The DAC settling requirement is the same as in the case of normal overranging: 1.4τ . The size of the next range is now $9/16$ of the previous range, resulting in 0.83 resolved bits per cycle and only seven instead of nine cycles are required for 6 bits of accuracy. Note that the DAC signal is no longer settling to the middle of the range, but instead to $9/16$ of the range, in order to give both new ranges (A and B) the same size.

For comparison, the diagram of an alternative architecture is shown in Fig. 13, where two comparators are used to select one out of three ranges. In this case, 1.4τ of settling is also required. The radix is 2, so six cycles are needed for 6 bits of accuracy. This architecture is commonly used in 1.5 bit/stage pipeline converters.

In Table I, an overview of the architectures is given. For each architecture the following information is shown: 1) the number of steps; 2) the required DAC settling time per step; 3) the total settling time; and 4) a number representing energy per conversion to compare the power efficiency. This number is proportional to the total conversion time and is doubled for the two-comparator case, where two comparators and two DACs are needed. This result is valid when static power conversion dominates over dynamic power consumption, which is true for our design.




Fig. 14. Photograph of the time-interleaved ADC and zoom-in on the T&H. The area of the T&H is 0.14 mm^2 and the total active area is 1.6 mm^2 .

From the table, it becomes clear that the architecture with two comparators has the shortest settling time and the single-sided overrange architecture is second best on this criterion, but that the latter has a much better power efficiency. Compared with the normal overrange architecture commonly used in SA-ADCs [8], the single-sided overrange architecture uses 25% less energy per conversion.

C. Look-Ahead Logic

To reduce the delay of the logic, we use look-ahead logic [8]: after each comparator decision there are two possible DAC levels, which are calculated in advance. Once the decision is known, the correct level only has to be picked and any logic gate delay is avoided.

V. CALIBRATIONS

Time-interleaved ADCs often require calibration of gain and offset [1], [2] and sometimes of timing [1]. Here, calibration of channel gain and offset is used to correct for deviations caused by the use of small T&H buffers [3]. This way no part of the input window is sacrificed and high-speed power-hungry digital operations are avoided. To enable the use of small differential pairs to keep the load for the DAC small and thus save power, comparator offset is made adjustable. All adjustments are controlled digitally by modifying analog bias settings in the input buffer via 6 and 7 bits DACs. In this test chip, the digital bias settings are controlled manually. The calibration could be automated in a start-up calibration, only requiring quasi-DC input signals.

The 16 times interleaved T&H and 16 connected ADCs are fabricated in $0.13 \mu\text{m}$ CMOS and a photograph of the chip is shown in Fig. 14, with a zoom-in on the interleaved T&H on the right-hand side, which has an area of 0.14 mm^2 . The total active area of T&H and ADCs is 1.6 mm^2 .

VI. MEASUREMENT RESULTS

First, the measurement result of a single channel is discussed. During this measurement all channels are active, however only the data from one channel is analyzed. The T&H is directly connected to a 50Ω signal generator. Our digital tester is limited to a few hundred megahertz, therefore we use on-chip decimation with a factor 9. In Fig. 15, the measurement result is shown at a total sample rate of 1350 MS/s resulting in $1350/16 \approx 84.4 \text{ MS/s}$ for a single channel. At low signal frequencies, the SNDR is

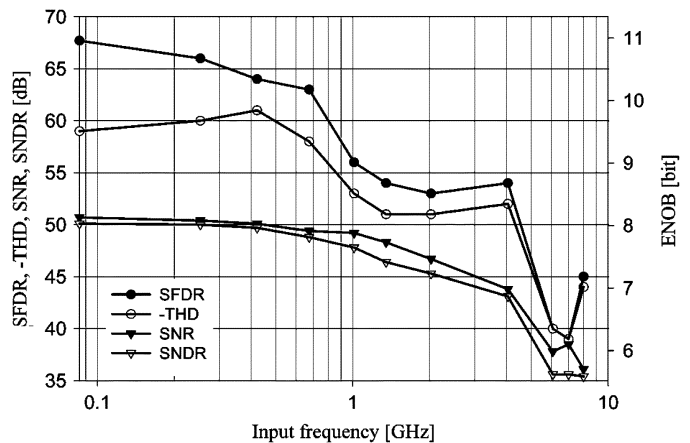


Fig. 15. Measurement result of a single channel with $f_s = 1350/16 \text{ MS/s} = 84.4 \text{ MS/s}$.

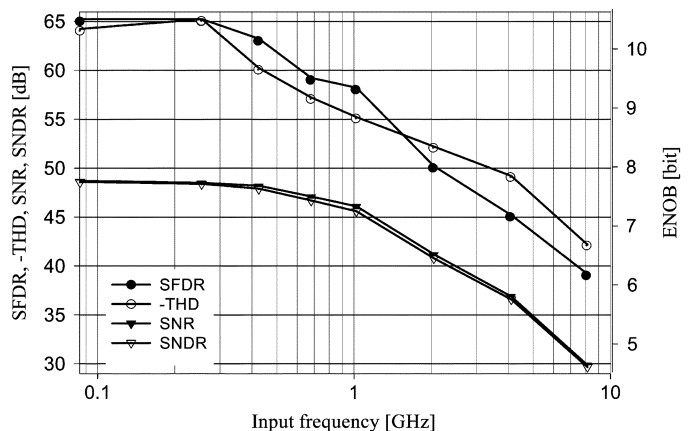


Fig. 16. Measurement result of the complete time-interleaved ADC at 1.35 GS/s .

8.0 ENOB limited by amplifier noise and quantization imperfections. The THD for low frequencies is -60 dB and the THD improvement at 8 GHz is due to a decrease in signal amplitude caused by losses in the test-bench signal path. THD at 4 GHz is -52 dB and THD at 8 GHz is -44 dB , which shows the excellent bandwidth and linearity of the T&H thanks to the use of the new circuit techniques applied in the T&H and the buffer. At 4 GHz input frequency, the SNDR is 6.9 ENOB and at 8 GHz the SNDR is 5.6 ENOB , which are higher than any values reported in literature for any ADC at these frequencies.

When increasing the signal frequency, thermal noise and quantization noise stay constant, while noise due to jitter increases linearly with the signal frequency. At a very high signal frequency, the SNR is strongly dominated by jitter and a good (but worst-case) approximation of the rms jitter is therefore given by: $\sigma(\Delta t) = 10^{-\text{SNR}/20} / 2\pi \cdot f_{\text{in}}$. Using this, the total jitter stemming from clock and signal generators and the circuit is only 0.2 ps rms , which is better than any value published for a T&H or ADC in CMOS.

The 16-channel interleaved performance at 1350 MS/s is shown in Fig. 16. The SNDR is 7.7 ENOB at low input frequencies and the ERBW is 1 GHz . Compared with the single channel case, the performance is only slightly degraded,




Fig. 17. Extracted timing offset (normalized).

showing that channel gain and offset are adjusted satisfactory and the step-size of the adjustment DACs is sufficiently small.

It is possible to extract the timing misalignment from the measured data by determining the phase of the output signal for each channel by means of an FFT. This way, jitter is averaged out and only the timing offsets remain. Note that bandwidth mismatch between channels also appears as timing offset, however, the expected amount is small compared with the timing offsets caused by sampling. The result of this operation for two measurements is shown in Fig. 17. The extracted rms timing misalignment is 0.6 ps rms, which is close to the expected value of 0.45 ps rms and which shows that the low skew technique is useful. Because of the dominance of timing misalignment, total timing error across all channels including jitter is also 0.6 ps rms. For ADCs with an ERBW over 500 MHz, this value is slightly better than the best reported in literature [1] where elaborate timing calibration is used. At 2 GHz, the SNDR is 6.5 ENOB and at 4 GHz the SNDR is 5.8 ENOB, limited by timing misalignment.

The input capacitance is about 1 pF and with a termination of 50 Ω on-chip and 50 Ω source impedance this results in an RC-limited analog input bandwidth of 6 GHz. The T&H buffers use a supply of 1.6 V, but all transistors have less than 1.2 V between their terminals. The rest of the circuit uses a supply voltage of 1.2 V.

Power consumption of the T&H including clock buffer and timing generation is 34 mW, the T&H buffers consume 40 mW and the 16 ADCs consume 100 mW. The FoM of the complete ADC calculated by $\text{power}/\{2^{\text{ENOB}} \cdot \min(f_s, 2 \cdot \text{ERBW})\}$ is 0.6 pJ per conversion step.

VII. IMPROVED DESIGN

In order to increase the performance of the ADC even more, two aspects are improved: 1) SNR and 2) sample rate. SNR is mainly limited by (1a) thermal noise of the inter-stage amplifier and (1b) DNL of the SA-ADCs. Noise from the inter-stage amplifier (1a) was decreased by circuit scaling. DNL (1b) was impaired by a parasitic capacitance asymmetry of 0.2 fF, causing crosstalk from the comparator to the DAC, resulting in an LSB error. Shielding or increasing wire distance was insufficient. Instead, the differential DAC outputs were twisted in the middle, making the crosstalk common mode and easy to reject by the comparator.

Significantly improved timing alignment and a higher maximum sample rate were achieved by boosting the bias current

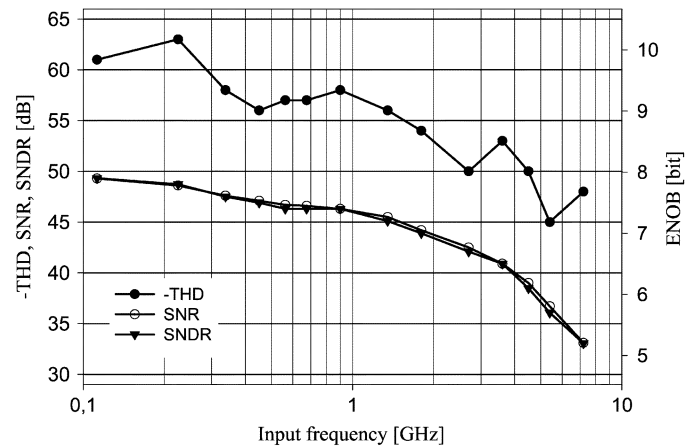


Fig. 18. All-channel measurement result of the improved design at 1.8 GS/s.

of the CML clock generator. Special care was taken with the dummy metal fill to avoid degradation of the maximum sample rate.

A. Measurement Results of Improved Design

At low sample rates, the interleaved performance is now 8.6 ENOB (8.8 ENOB for a single channel) proving that the noise of the inter-stage amplifier is lowered and the DNL of SA-ADCs is reduced.

At the nominal supply voltage, the T&Hs and SA-ADCs are functional up to 2 GS/s, however, the inter-stage amplifier is only functional up to a sample rate of 1.2 GS/s, for higher frequencies the differential output is zero. When the bias current of the amplifier is decreased, it is functional at higher sample rates, however in this case its limited settling degrades the SNDR. In order to make the amplifier operate with nominal bias settings, the supply voltage has to be increased to 1.8 V, which degrades the SA-ADCs performance. At a sample rate of 1.8 GS/s, single-channel performance is 8.3 ENOB for low input frequencies, 7.4 ENOB@3.6 GHz and 5.9 ENOB@7.2 GHz.

Measurement results using all channels at 1.8 GS/s are shown in Fig. 18. At low input frequencies, the SNDR is 7.9 ENOB, limited by DAC settling and the negative effect of the high supply voltage. The ERBW is 1 GHz and the FoM is 1 pJ/conversion step and the power consumption is 416 mW, which is almost doubled due to the increased supply voltage. Total timing error due to jitter and misalignment between channels is 0.4 ps rms. The timing alignment is improved by 30% due to the increased edge steepness of the CML clock.

Compared with the first design the maximum sample-rate is increased significantly and at nominal supply (and reduced sample rate) the SNDR is increased by almost 1 ENOB. Table II shows an overview of state-of-the-art time-interleaved ADCs. Reference [1] has a very high sample rate but is not suitable for embedded application and requires timing calibration. Our design reaches the same timing alignment without timing calibration and has less jitter. Compared with [2] and [6], we present a much higher sample rate and ERBW. Beyond $f_{in} = 1$ GHz, our design achieves better performance than ever published before.

TABLE II
PERFORMANCE OVERVIEW OF HIGH-SPEED TIME-INTERLEAVED ADCs

Design	Poulton 2003 [1]	Gupta 2006 [2]	HSU 2007 [6]	Louwsma 2007 [5,9]	Improved design
Sample-rate [GS/s]	20	1.0	0.8	1.35	1.8
ENOB ($f_{in} = \text{DC}$)	6.5	8.85	9.0	7.7	7.9 (1-ch: 8.3)
ENOB ($f_{in} = 4 \text{ GHz}$)	5.3	-	-	5.8	6.4 (1-ch: 7.3)
ERBW [GHz]	2.0	0.4	0.4	1.0	1.0
Input bandwidth [GHz]	6.6	-	-	6	6
Power consumption [W]	10	0.25	0.35	0.18	0.42
$\text{FoM} = \frac{P}{2^{\text{ENOB}} \cdot \min(f_s, 2 \cdot \text{ERBW})}$ [pJ/conv.step]	28	0.7	0.9	0.6	1
Jitter [ps RMS]	0.6	-	0.43	0.2	0.2
Timing misalignment [ps RMS]	0.4 (calibrated)	-	-	0.6	0.4

VIII. CONCLUSIONS

A time-interleaved ADC is demonstrated with 16 channels. By using a new buffer and removing the load in track mode, the T&H reaches a high bandwidth and good linearity. For a single channel, THD is -52 dB at an input frequency of 4 GHz and SNDR is 43 dB, which is only limited by (best-in-class) jitter of 0.2 ps rms. With a novel circuit design a good timing alignment of 0.6 ps rms is achieved, even without timing calibration.

By pipelining two SA-ADCs, a combination of high sample rate and good power efficiency can be reached. The single-sided overrange architecture achieves a 25% higher power efficiency compared with the conventional overrange architecture and look-ahead logic minimizes logic delay in the SA-ADC. The FoM of the complete ADC including T&H is 0.6 pJ per conversion step. The SNDR is 7.7 ENOB for low signal frequencies, while the ERBW is 1 GHz, showing broadband signal handling capability.

An improved design achieves an SNDR of 8.6 ENOB for low sample-rates and with a higher supply voltage it reaches a sample rate of 1.8 GS/s with 7.9 ENOB at low signal frequencies and an ERBW of 1 GHz. At $f_{in} = 3.6$ GHz, the SNDR is still 6.5 ENOB and total timing error including jitter is only 0.4 ps rms, which is better than any value published for an ADC with a bandwidth larger than 500 MHz.

ACKNOWLEDGMENT

The authors would like to thank H. v.d. Ploeg, G. Wienk, J. Briaire, and K. Doris for their useful discussions and assistance and NXP Semiconductors for fabrication of the chip.

REFERENCES

- [1] K. Poulton *et al.*, "A 20GS/s 8b ADC with a 1MB memory in 0.18 μm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2003, pp. 318–319.
- [2] S. Gupta, M. Choi, M. Inerfield, and J. Wang, "A 1GS/s 11b time-interleaved ADC in 0.13 μm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2006, pp. 264–265.
- [3] S. M. Louwsma, E. J. M. van Tuijl, M. Vertregt, P. C. S. Scholtens, and B. Nauta, "A 1.6 GS/s, 16 times interleaved track & hold with 7.6 ENOB in 0.12 μm CMOS," in *Proc. ESSCIRC*, Sep. 2004, pp. 343–346.

- [4] K. Nagaraj *et al.*, "A dual-mode 700-Msamples/s 6-bit 200-Msamples/s 7-bit A/D converter in a 0.25 μm digital CMOS process," *IEEE J. Solid-State Circuits*, vol. 35, no. 12, pp. 1760–1768, Dec. 2000.
- [5] S. M. Louwsma, E. J. M. van Tuijl, M. Vertregt, and B. Nauta, "A 1.35 GS/s, 10b, 175 mW time-interleaved AD converter in 0.13 μm CMOS," in *Symp. VLSI Circuits Dig.*, Jun. 2007, pp. 62–63.
- [6] C.-C. Hsu *et al.*, "An 11b 800MS/s time-interleaved ADC with digital background calibration," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 464–465.
- [7] W. C. Black and D. A. Hodges, "Time interleaved converter arrays," *IEEE J. Solid-State Circuits*, vol. SC-15, no. 6, pp. 1022–1028, Dec. 1980.
- [8] F. Kuttner, "A 1.2V 10b 20MSample/s non-binary successive approximation ADC in 0.13 μm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2002, pp. 176–177.
- [9] S. M. Louwsma, E. J. M. van Tuijl, M. Vertregt, and B. Nauta, "A time-interleaved track & hold in 0.13 μm CMOS sub-sampling a 4 GHz signal with 43 dB SNDR," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC)*, Sep. 2007, pp. 329–332.
- [10] N. Kurosawa, H. Kobayashi, K. Maruyama, H. Sugawara, and K. Kobayashi, "Explicit analysis of channel mismatch effects in time-interleaved ADC systems," *IEEE Trans. Circuits Syst. I, Fundam. Theory Appl.*, vol. 48, no. 3, pp. 261–271, Mar. 2001.
- [11] A. M. Abo and P. R. Gray, "A 1.5-V, 10-bit, 14.3-MS/s CMOS pipeline analog-to-digital converter," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 599–606, May 1999.



Simon M. Louwsma was born on January 1, 1976, in Wommels, The Netherlands. He received the M.Sc. degree in electrical engineering from the University of Twente, Enschede, The Netherlands, in 2001, where he is currently working toward the Ph.D. degree in the IC Design Group on time-interleaved ADCs.

He holds several patents and is cofounder of Axiom-IC, which specializes in data-converters. His interests include high-speed, low-power ADCs and other analog and mixed-signal circuits and systems.



A. J. M. van Tuijl (M'85) received the degree in electrical engineering from the Delft University of Technology, Delft, The Netherlands, in 1978.

From 1980 to 1998, he was with Philips Semiconductors, Nijmegen, The Netherlands, where he worked on many different mixed-signal applications as a Designer and as a Design Manager. In 1992, he joined the University of Twente, Enschede, The Netherlands, as a part-time Professor of analog electronics. In 1998, he joined Philips Research Eindhoven, Eindhoven, The Netherlands, as a Principal Research Scientist in the field of high-speed data conversion and wired data Communication. He is the author or coauthor of many papers and holds many patents in the field of analog electronics and data conversion. He was recently part of the start-up of a new company, Axiom-IC, Twente, The Netherlands, that specializes in custom design and IP delivery of very high-speed data converters. The new company has strong links with IC Design Group of the University of Twente, where he is still a part-time Professor.



Maarten Vertregt received the M.Sc. degree in electrical engineering from the University of Twente (UT), Enschede, The Netherlands.

He joined Philips Research in 1985. He worked on the design of 1 Mbit and 4 Mbit SRAM memories and subsequently on A/D conversion with embedded signal processing. Since 1996, he has coordinated the design activities for high-speed A/D and D/A conversion functions within the Mixed-Signal Circuits and Systems Group of Philips Research. Presently, he is a Senior Principal with the Research Department (Corporate Innovation and Technology), NXP Semiconductors, Eindhoven, The Netherlands. His research interests are with the migration of signal processing from the analog domain to the digital domain to cope with new system demands,

in combination with scaling to nanometer-CMOS process technology. He was a plenary invited speaker on analog scalability at the 2004 ESSCIRC/ESSDERC and 2006 IEDM conferences. He has authored and coauthored over 20 publications and nine patents.



Bram Nauta (M'91–SM'03) was born in Hengelo, The Netherlands, in 1964. He received the M.Sc. degree (*cum laude*) in electrical engineering and the Ph.D. degree from the University of Twente, Enschede, The Netherlands, in 1987 and 1991, respectively. His doctoral work focused on analog CMOS filters for very high frequencies.

In 1991, he joined the Mixed-Signal Circuits and Systems Department of Philips Research, Eindhoven, The Netherlands, where he worked on high-speed A/D converters and analog key modules.

In 1998, he returned to the University of Twente as full Professor heading the IC Design group, which is part of the CTIT Research Institute. His current research interest is high-speed analog CMOS circuits. He is also a part-time consultant in industry.

Prof. Nauta served as Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS PART II—ANALOG AND DIGITAL SIGNAL PROCESSING (1997–1999), and as Guest Editor (1998) and Associate Editor (2001–2006) for the IEEE JOURNAL OF SOLID-STATE CIRCUITS. He is a member of the technical program committees of ISSCC, ESSCIRC, and the Symposium on VLSI Circuits. He was corecipient of the ISSCC 2002 Van Vessel Outstanding Paper Award. He is currently Editor-in-Chief of the IEEE JOURNAL OF SOLID-STATE CIRCUITS.