A 12-bit 2-GS/s TI Pipe-SAR ADC with Optimized Signal Distribution and Fully On-Chip Calibrations

Alessia Ceroni*,1, Michele Rocco*,1,2, Giacomo Tombolan*,1, Gabriele Zanoletti[1], Gabriele Bè[1,2], Luca Ricci[1], Salvatore Levantino[1], Andrea Leonardo Lacaita[1], Luca Bertulessi[1], Carlo Samori[1], Andrea Giovanni Bonfanti[1]
[1]*Politecnico di Milano*, [2]*now with Infineon Technologies, Villach, Austria*, **Equally Credited Authors (ECAs)*

***Abstract*—This paper presents a 2-GS/s time-interleaved (TI) ADC based on four pipeline successive approximation register (pipe-SAR) slices and addresses system-level challenges arising in TI operation. An optimized signal-distribution network is proposed, featuring an input buffer and a symmetric Y-tree that minimizes the clock–input coupling without requiring shielding or intertwisting, thus preserving the signal bandwidth. Fully on-chip calibrations correct offset, gain, and timing-skew mismatches. The skew calibration is combined with a monotonic digital-to-time converter (DTC) control strategy to ensure fast convergence. Fabricated in 28-nm CMOS bulk technology, the prototype dissipates 87.3 mW—including the input buffer and on-chip digital calibrations—and achieves a 61.6/79.2-dB SNDR/SFDR near Nyquist. Compared to the single-core operation, the average SNDR/SFDR performance drop is only 0.5/0.4 dB over the whole signal bandwidth.**

***Index Terms*—TI ADCs, Input Buffer, Signal Distribution Tree, TI calibrations**

I. Introduction

Analog-to-digital converters (ADCs) targeting large bandwidths at medium-to-high resolutions commonly rely on time-interleaving [1]. However, this approach introduces critical challenges, such as front-end linearity degradation at high frequencies, clock skew, and channel mismatches [2]. Mitigating these non-idealities requires additional circuitry and significant power overhead. To minimize the interleaving factor—and thereby reduce the impact of TI-related issues—the pipe-SAR architecture represents a preferred choice to implement the slice, as it requires fewer channels for a target bandwidth [3].

This work presents a 2-GS/s 4× TI pipe-SAR ADC (Fig. 1) with an optimized signal distribution network and fully on-chip calibrations, achieving a 61.6/79.2-dB SNDR/SFDR at a 1-GHz input. The overall power consumption is 87.3 mW, resulting in a Schreier's figure of merit (FoMs) of 162.2 dB.

II. Time-Interleaved ADC Implementation

The differential input and clock signals are terminated on-chip, buffered near the pads, and distributed via a symmetric Y-tree to the 4 sub-ADCs. The slices use the architecture in [4]. Their outputs are aligned to a common clock and processed in the digital back-end. Background calibration corrects comparator offset, interstage gain error, and interleaving mismatches, while the CDAC radix errors are handled in the foreground.

A. Input buffer

An on-chip input buffer drives the ADC, providing isolation. An open-loop push-pull topology (Fig. 2a) is adopted

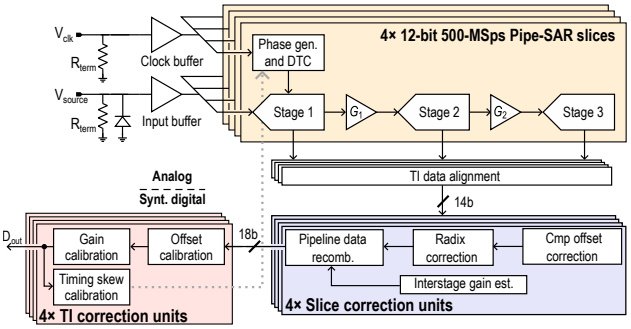

Fig. 1: Block diagram of the 4× TI pipe-SAR ADC.

to achieve the 1-GHz bandwidth requirement with minimal power consumption. To enhance the linearity, the output stage employs two-level drain bootstrapping [5]. The input network features on-chip 50-Ω termination resistors (R_{term}) for impedance matching and custom ESD protections. The signal is AC-coupled to the push-pull and cascode gates via 750-fF capacitors (C_{LS}) and isolated from the common-mode feedback (CMFB) circuit by 200-kΩ resistors (R_{LS}). The output stage uses minimum-length core transistors rather than thick-oxide devices: The reduced parasitics and improved output headroom maximize bandwidth and linearity. The reliability has been verified via extensive simulations.

Figure 2b shows the simulated HD3 of the implemented buffer versus the input frequency (blue curve). At low frequencies, the HD3 is dominated by the drain current nonlinearity. Around 100 MHz, the distortion increases due to the nonlinear input capacitances, primarily those of the ESD devices and the gate-drain capacitances (C_{gd}) of the push-pull pair. To mitigate this effect, the ESD devices are sized to the minimum value that ensures adequate protection. Compared to the green curve with standard protection devices, the HD3 is improved by almost 10 dB in the mid-frequency range. Additionally, the two-level drain-bootstrapped output stage maintains nearly constant V_{DS}. This significantly reduces the C_{gd} modulation compared to single-level bootstrapping (purple curve), which suffers from residual V_{DS} variation as the bootstrap device drains remain at fixed potentials [5]. Stable V_{DS} and C_{gd} enhance the HD3 by almost 20 dB around 100 MHz. At high frequencies, the signal-dependent load current causes a -60 dB/dec HD3 degradation. A peak near 3 GHz results from a partial cancellation between the nonlinear input capacitance and load current distortion mechanisms.

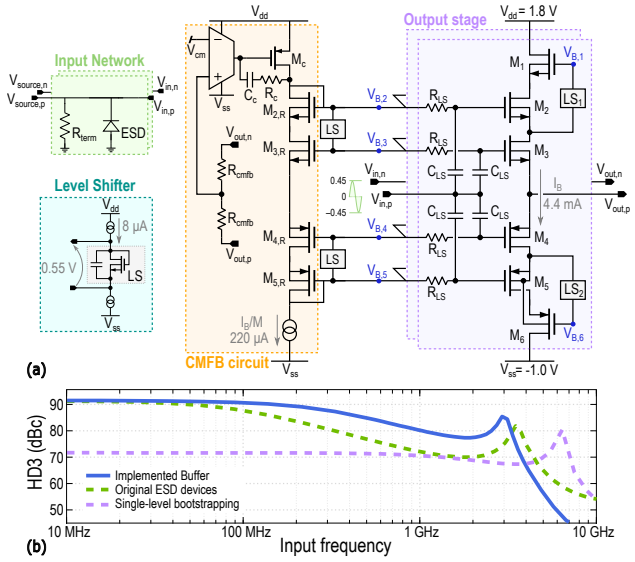

Fig. 2: Input buffer (a). Simulated HD3 magnitude (b).

The CMFB circuit senses the output common mode through 100-kΩ polysilicon resistors (R_{cmfb}) and generates (i) the cascode bias voltages $V_{B,3-4}$ to regulate the output at V_{cm} = 0.45 and (ii) $V_{B,2-5}$ to properly bias the cascode devices M_{2-5}. The CMFB employs an operational transconductance amplifier (OTA) and a matched diode-connected replica of the output stage scaled by $M = 20$ and biased at 220 μA. Level shifters (LS_{1-2}) with diode-connected transistors provide the DC bias for the stacked devices M_{1-6} ($V_{B,1-6}$), offering superior process tracking over resistive level shifters.

B. Input and clock distribution network

In the clock and input distribution network, the coupling between these signals is a major concern as it degrades the performance. In TI ADCs, this issue is exacerbated: Both signals are routed in parallel over long distances to the slices, significantly increasing the coupling. Let $\alpha(f)$ and $\beta(f)$ denote the input-to-clock and clock-to-input coupling transfer functions, respectively, and consider the differential input $V_{in}(t) = A_{in}\sin(2\pi f_{in}t + \phi_{in})$ and clock $V_{clk}(t) = A_{clk}\sin(2\pi f_s t)$ signals. The input-to-clock coupling introduces a signal-dependent perturbation on the sampling instant ($T_s = 1/f_s$), causing sampling time modulation ($T_s \rightarrow T_s + \Delta T_s$). Its effect on the sampled input can be assessed as

$$V_{in}(nT_s + \Delta T_s) \approx V_{in}(nT_s) + \Delta T_s \frac{dV_{in}}{dt}, \qquad (1)$$

with $\Delta T_s(t)$ being proportional to $\alpha(f_{in})$ and the input signal. Both $\Delta T_s(t)$ and dV_{in}/dt depend on the input frequency, thus their product generates a second harmonic, namely,

$$\text{HD2} \propto \frac{f_s A_{clk}}{|\alpha(f_{in})| f_{in} A_{in}}. \qquad (2)$$

The clock-to-input coupling, instead, produces an additive disturbance proportional to $|\beta(f_s)|$ on V_{in}, resulting in a DC offset when sampled at the clock instants. The interleaving offset calibration can correct this error, and its digital implementation can easily handle the additional correction range.

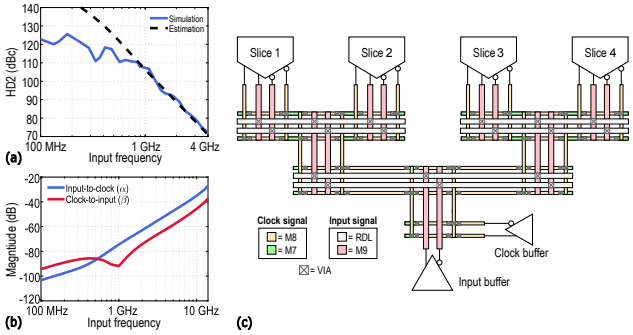

Fig. 3: Estimated and simulated HD2 (a). In-to-clk and clk-to-in transfer functions from the simulations of the front-end (input buffer, EM-extracted tree, sampler) (b). Proposed tree structure (c).

Coupling is conventionally mitigated through ground shielding or intertwisted routing. Shielding reduces the crosstalk by inserting grounded traces between the clock and input lines [6], but it increases the parasitic capacitance and thus limits the bandwidth [7]. Intertwisting achieves a first-order coupling cancellation through periodic alternation of the signal segments [7], but it introduces impedance imbalance and exacerbates the timing skew and bandwidth mismatches.

In this work, the intrinsic coupling of a balanced distribution network (without shielding or inter-twisting) is investigated not to limit the HD2 across the 1-GHz bandwidth. A tree topology is chosen to match all the clock and input signal paths to the four channels, avoiding impedance unbalances, minimizing skew and bandwidth mismatches, without excessive loading on the input nodes. The transfer function $\alpha(f)$ is extracted through electromagnetic (EM) simulations for various input-clock lines spacings and then used, together with the simulated clock slope (100 mV/ps), to analytically estimate the HD2 from Eq. (2). The estimated HD2 (Fig. 3a) decreases with f_{in} at a rate of approximately -60 dB/dec, which is consistent with the -40 dB/dec roll-off of $\alpha(f_{in})$ (Fig. 3b) and the $1/f_{in}$ term in Eq. (2). Figure 3a demonstrates a close agreement for the HD2 between the EM simulations and the analytical model at high frequencies, where this effect is the dominant one. The adopted distribution network (Fig. 3c) employs different thick metal-layer pairs for better isolation: The input is routed with 3.6-μm-wide traces on the highest metal level (M9) and redistribution layer (RDL), while the clock uses 1.8-μm-wide traces on M7 and M8. The signal and clock lines are spaced by 3.6 μm, and the total height is 122 μm, resulting in a compact form factor. The design achieves an HD2 $>$ 100 dB at a 1 GHz input frequency, confirming that the coupling does not limit the linearity within the target bandwidth.

C. On-chip interleaving calibrations

Interleaving calibrations correct the offset by subtracting the time-average of each sub-ADC output and equalize the gain across the slices with respect to a reference channel. The skew calibration uses a mixed-signal approach: The skew is

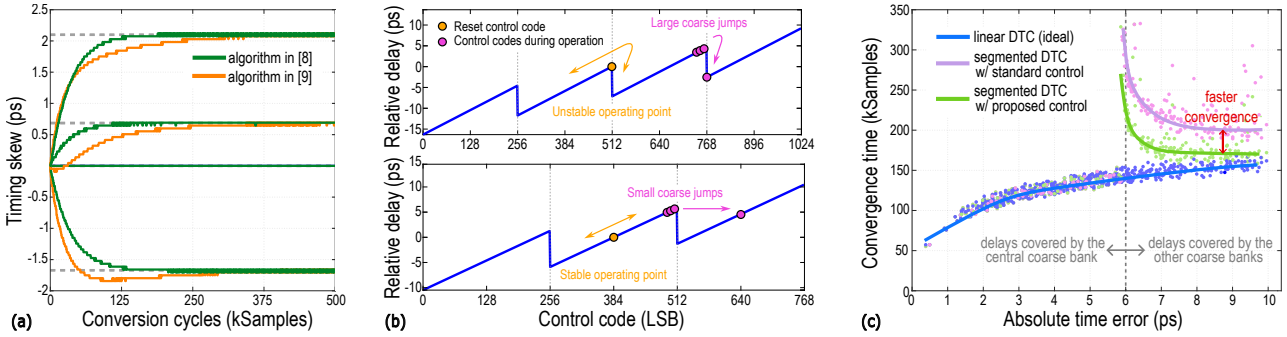

Fig. 4: Simulated skew calibration transient based on [8] (green) and [9] (orange) (a). Standard vs proposed DTC control (b). Skew calibration convergence time for the linear, segmented with standard control, and segmented with proposed control DTCs. With segmented DTCs, the convergence time increases when the skew error is larger than the maximum delay covered by the central coarse bank (c).

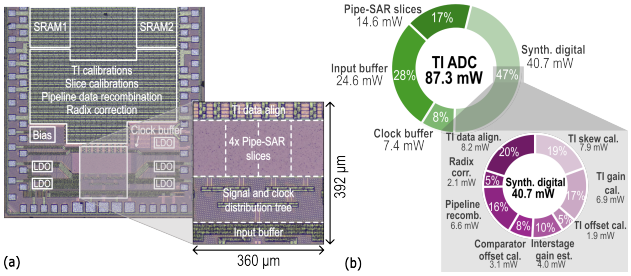

Fig. 5: Die micrograph (a). Measured power breakdown (b).

estimated digitally from cross-correlations of the channels' outputs and corrected by adjusting the sampling instants through DTCs. For the i^{th} channel, reference [8] uses

$$e_i = M\sum_{j=0}^{i-1} r_{j,j+1} - (i-1)\cdot\sum_{i=0}^{M-1} r_{i,i+1} \quad \propto (\tau_i - \tau_0) \quad (3)$$

as the error signal to estimate the skew. Here, $r_{a,b}$ is the correlation between channels a and b, M is the interleaving factor, and τ_i, τ_0 are the skew of the i^{th} and reference channels. This approach is advantageous compared to typical correlation-based techniques [9] since the error signal depends only on the channel under calibration. In contrast, in the latter, the skew is estimated in parallel by cross-correlating adjacent channels, leading to coupled error signals and slower convergence due to overshoots/undershoots [8]. Thus, the implemented skew calibration algorithm is based on [8], but here the Mean Absolute Difference (MAD) is replaced with actual correlations to avoid a biased convergence in multi-tone scenarios. Figure 4a compares the convergence of both methods when the same loop bandwidth[1] is used, clearly showing overshoots for [9] (orange curve). With convergence defined as the DTC control settling within ± 2 LSB around the steady-state value, the orange curve converges in 300 kS, while the proposed approach (green curve) requires only 125 kS.

The convergence also depends on the correction actuator. The DTC is implemented by inverters loaded with digitally controllable metal-oxide-metal (MOM) capacitive banks. It is partitioned into a 3 thermometric-level coarse and an 8-bit fine stage for a low jitter-power product [10]. Segmented DTCs are typically designed with overlapping coarse steps to avoid missing codes in the delay characteristic due to PVT variations. Here, the nominal coarse LSB (5 ps) is smaller than the fine range in the fastest corner (5.3 ps), resulting in a large overlap in the typical case. Thus, when transitioning to the next coarse bank, the effective delay significantly decreases instead of slightly increasing, resulting in a larger convergence time (Fig. 4b, top). To avoid this, when the DTC control word overflows/underflows, the fine DTC control word is set to the mid-range of the fine bank rather than to its minimum/maximum value, yielding a delay closer to the maximum/minimum of the previous coarse bank (Fig. 4b, bottom). This approach preserves the delay monotonicity and significantly accelerates the convergence. Figure 4c compares the convergence time for three DTC implementations: ideally linear (blue), segmented with standard control (purple), and segmented with monotonic control (green). A Monte Carlo simulation generates the skew distribution for the channels. For each simulation point, the calibration algorithm is executed, and the convergence times of the three implementations are recorded against the maximum absolute skew among the channels. Solid lines interpolate the simulation points, revealing that the proposed monotonic control achieves faster convergence. Lastly, the proposed DTC uses an odd number of coarse levels such that the reset control code corresponds to the center of a fine delay characteristic, providing a stable operating point at startup (Fig. 4b, bottom).

In the on-chip digital implementation, the most area- and power-demanding blocks are the correlators (multipliers). To minimize the power consumption, the hardware is maximally reused: Each skew-calibration engine computes only the correlation between its assigned channel and the following one. By sharing the results across the four engines, only one multiplier per engine is needed (four total).

III. Measurement Results

The TI ADC is fabricated in a TSMC 28-nm bulk CMOS technology and occupies an active area of 0.141 mm^2 (Fig. 5a). The core and reference buffer supplies are 1 V and 1.2 V, respectively; the input buffer is supplied between 1.8 V and

[1]The error signal in [8] is down-scaled by $M/2$ with respect to [9] to guarantee that the same error is integrated in the calibration loop.

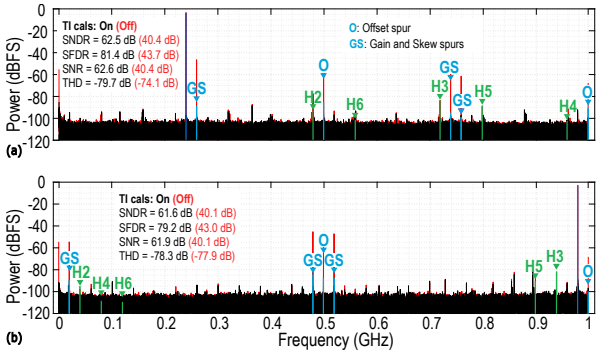

Fig. 6: Measured output spectra at 240 MHz (a) and 980 MHz (b).

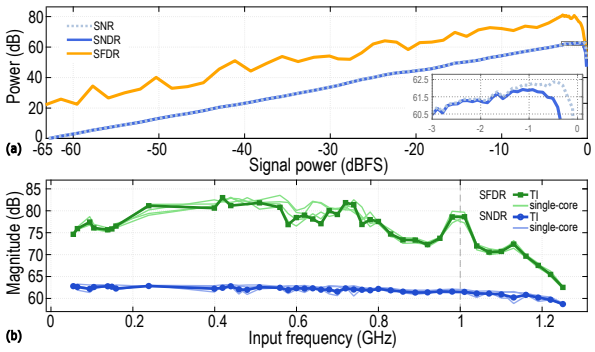

Fig. 7: Measured input amplitude (a) and frequency (b) sweeps.

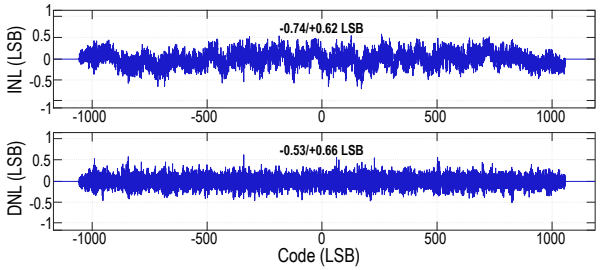

Fig. 8: Measured INL (top) and DNL (bottom).

	This work		[11]	[12]	[13]	[14]
ADC architecture	TI Pipe-SAR		TI Pipe	Pipe	Pipe+TI SAR	Pipe+TI SAR
Input freq. (MHz)	980		1600	991.5	896.48	988
SNDR (dB)	61.6		61.7	60.4	60.2	58.8
SFDR (dBc)	79.2		73.3	75.8	69.2	70.7
Sampl. freq. (GHz)	2		3.2	2	1.8	3
Power (mW)	46.6*	87.3†	61.3	27	7.55	32.5
Walden FoM (fJ/c.s.)	23.7	44.4	19.3	15.8	5.0	15.2
Schreier FoM (dB)	164.9	162.2	165.9	166.1	171.0	165.4
Active area (mm²)	0.141		0.194	0.034	0.0243	0.04
Technology (nm)	28		16	28	7	28
Slice calibrations	BG-on (gain, offset) FG (radix)		BG-off (RA NL) FG-off (gain)	FG-off (gain, offset, radix)	FG-off (radix)	BG-on (RA NL) FG-off (gain, offset)
TI calibrations	BG-on (offset, gain, skew)		FG-off (skew)		No	
Input buffer	Yes		Yes	No	No	No

*: analog only (pipe-SAR slices, Input Buffer, Clock Buffer), †: analog and digital calibrations

TABLE I: Performance summary and comparison with the SotA.

-1 V. The total power consumption is 87.3 mW, including the digital calibrations and the input buffer (Fig. 5b).

Figure 6 shows the output spectra for 240-MHz and 980-MHz input tones. The measured SNDR/SFDR are 62.5/81.4 dB at low frequency and 61.6/79.2 dB near Nyquist. The interleaving spurs are suppressed below -80 dBFS: The SNDR is limited by the thermal noise, whereas the SFDR is limited by the third harmonic. Figure 7a shows the amplitude sweep for an input tone at 980 MHz, while Fig. 7b reports the input frequency sweep with an amplitude of -1.45 dBFS. Here, the blue and green dark lines show the SNDR and SFDR of the TI ADC, respectively, while the light ones refer to the four sub-ADCs. Over the entire frequency range, the average SNDR/SFDR spread between TI and single-core operation is only 0.5/0.4 dB. With all the calibrations active, the measured INL and DNL with a Nyquist input are between $-0.74/+0.62$ LSB and $-0.53/+0.66$ LSB, respectively (Fig. 8). The FoMs is 164.9 dB when accounting only for the analog power, and 162.2 dB when including also the digital. Table I summarizes the measured performance and compares it with state-of-the-art ADCs with a similar bandwidth. This work reports the highest SFDR among the compared designs.

IV. Conclusions

This work presents a 2-GS/s 12-bit TI pipe-SAR ADC in 28-nm CMOS. The two-level drain-bootstrapped input buffer with custom ESD protection guarantees sufficient linearity, and the EM-optimized signal distribution network suppresses the signal coupling, without requiring complex isolation techniques. The skew calibration combines the algorithm from [8] (implemented with actual correlations) with a monotonic DTC control strategy, eliminating delay backtracking and accelerating the convergence speed. The design achieves a 61.6/79.2-dB SNDR/SFDR near Nyquist and has a minimal performance degradation compared to the single-core operation, validating the proposed TI-mitigation techniques.

References

[1] Black, W., et al. "Time interleaved converter arrays." IEEE JSSC (2003).
[2] Ricci, L., et al. "A 2GS/s time-interleaved ADC with embedded background calibrations and a novel reference buffer for reduced inter-channel crosstalk." IEEE JSSC (2024).
[3] Lee, C., et al. "A SAR-assisted two-stage pipeline ADC." IEEE JSSC (2011).
[4] Rocco, M., et al. "A 500MS/s 12b Pipe-SAR ADC Using a Triple-Cascode FIA with Virtual Supply Extension." IEEE ISSCC (2026).
[5] Devarajan, S., et al. "A 12b 10GS/s interleaved pipeline ADC in 28nm CMOS technology." IEEE JSSC (2017).
[6] Nam, J.-W., et al. "A 12bit 1.6, 3.2, and 6.4GS/s 4b/cycle time-interleaved SAR ADC with dual reference shifting and interpolation." IEEE JSSC (2018).
[7] Ramkaj, A., et al. "A 5GS/s 158.6mW 9.4ENOB passive-sampling time-interleaved three-stage pipelined-SAR ADC with analog–digital corrections in 28nm CMOS." IEEE JSSC (2020).
[8] Gu, M., et al. "A 1GS/s 11b time-interleaved SAR ADC with robust, fast, and accurate autocorrelation-based background timing-skew calibration." IEEE JSSC (2024).
[9] Bè, G., et al. "A 900MS/s SAR-based time-interleaved ADC with a fully programmable interleaving factor and on-chip scalable background calibrations." IEEE TCAS-II (2022).
[10] Santiccioli, A., et al. "Power-jitter trade-off analysis in digital-to-time converters." Electronics Letters (2017).
[11] Hershberg, B., et al. "A 3.2GS/s 10ENOB 61mW Ringamp ADC in 16nm with Background Monitoring of Distortion." IEEE ISSCC (2019).
[12] Cao, Y., et al. "A single-channel 12b 2GS/s PVT-robust pipelined ADC with critically damped ring amplifier and time-domain quantizer." IEEE ISSCC (2023).
[13] Hsieh, S., et al. "A 1.8GHz 12b pre-sampling pipelined ADC with reference buffer and OP power relaxations." IEEE ISSCC (2023).
[14] Gu, M., et al. "A 12b 3GS/s pipelined ADC with gated-LMS-based piecewise-linear nonlinearity calibration." IEEE ISSCC (2025).