22RPT02 True8DIGIT

D1: Report identifying novel metrology grade ADC architectures for the DC to 100 kHz frequency band. This will include comprehensive digital models covering integrating ADC (IADC) and a mixed design e.g. IADC with SAR, as well as first and second order error mechanisms.

Organisation name of the lead participant for the deliverable:
Left Right

Organisation name of the contributing participants for the deliverable:
CERN, CMI, FER, INTI, JanasCard, SC

Due date of the deliverable: 31. 1. 2026

Actual submission date of the deliverable: 28. 1. 2026

Confidentiality Status: PU –Public, fully open.

Deliverable Cover Sheet

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or EURAMET. Neither the European Union nor the granting authority can be held responsible for them.

The project has received funding from the European Partnership on Metrology, co-financed from the European Union's Horizon Europe Research and Innovation Programme and by the Participating States.

Table of Contents

1 Summary

This report addresses the feasibility of achieving true 8-digit (relative accuracy at the 10^{-8} level) digitisation in the DC to 100 kHz frequency range. The work focuses on identifying fundamental performance barriers in existing metrology-grade digitisers and evaluating architectural and signal-processing strategies capable of overcoming these limitations.

Three dominant barriers are identified: noise, linearity, and timing accuracy. Noise remains the primary limitation at room temperature, with white noise, $1/f$ noise, and long-term random walk each imposing different constraints depending on measurement bandwidth and aperture time. While integrating ADCs (IADCs) provide excellent suppression of white noise through time averaging, they are increasingly affected by low-frequency noise and timing-related effects at long integration times. Conversely, SAR-based ADCs offer superior timing determinism and bandwidth but are limited by achievable linearity and noise. These complementary characteristics motivate the exploration of alternative architectures that combine the advantages of both approaches.

A key contribution of this work is the introduction and analysis of a novel digitiser topology that combines integrating and sampling principles in a structured and scalable manner. The proposed architecture exploits the noise-rejection capability of integration while retaining precise timing control and high linearity through parallelisation and digital post-processing. By distributing the measurement task across multiple channels and combining their outputs in the digital domain, the topology offers a path to averaging down noise, mitigating nonlinearity, and relaxing individual component requirements. This approach is particularly well suited to wideband precision measurements where conventional single-channel architectures reach practical limits.

To support architectural development, extensive modelling was carried out at multiple levels. System-level models were used to study noise propagation, averaging behaviour, and long-term stability limits, while circuit-level simulations provided insight into non-ideal effects such as switching transients, charge injection, and finite amplifier bandwidth. Digital models complemented these efforts by evaluating post-processing strategies, including averaging, demodulation, and error separation. The combination of these modelling approaches enables consistent cross-verification of assumptions and highlights which error mechanisms are fundamental and which can be mitigated by design.

As part of the circuit-level investigation, switching-induced charge coupling in integrating stages was analysed using time-resolved LTspice simulations. While individual switching events produce fast current spikes, the simulations show that the net residual charge per complete switching cycle converges to a small and repeatable value. This indicates that switching-related effects behave predominantly as deterministic offsets rather than as stochastic noise sources, and therefore do not constitute a fundamental accuracy barrier when appropriate calibration or correction is applied.

In conclusion, the results presented in this report support the technical feasibility of developing a true 8-digit digitiser for DC with state-of-the-art stability and linearity up to 100 kHz operation. The dominant limitations are not architectural in nature but arise from practical implementation aspects such as parasitics, thermal stability, and long-term drift. The proposed digitiser topology, supported by multi-level modelling, provides a flexible foundation for overcoming these challenges. The logical next step is the development and experimental characterisation of a physical prototype to validate the predicted performance and to quantify real-world effects that cannot be fully captured in simulation.

We acknowledge the valuable contributions of our collaborator, Jaromir Sukuba (Slovakia), for his many insightful discussions that significantly informed this work.

2 Introduction

This report identifies novel metrology grade ADC architectures for the DC to 100 kHz frequency band. It covers particular performance barrier mitigations for Integrating ADC (IADC) and Successive Approximation Register (SAR) based analogue to digital converters, as well as comprehensive digital models covering IADC architectures. Finally, it proposes a hybrid operation of IADC and SAR ADCs and system wide mitigations.

The work builds upon the barriers identified in currently available digitisers, specifically noise, nonlinearity, and timing accuracy, and proposes circuit level and post-processing workarounds to reduce critical component requirements.

3 Barriers to digitiser performance

3.1 Noise

Current state-of-the-art digitisers face several fundamental barriers that limit their performance in the DC to 100 kHz range. As identified in project activity A1.1.3, these barriers primarily include noise, linearity, and timing accuracy.

Noise is an inherent limitation in any electronic measurement device, representing unwanted random fluctuations that obscure the true signal. For room-temperature digitisers, it acts as a fundamental barrier to precision. Every stage of a digitiser — from the initial input amplifier to the final analog-to-digital converter — adds its own noise to the measurement.

The most critical types of noise affecting these systems are:

1. **White (Gaussian) Noise:** This is a random background noise that is evenly distributed across all frequencies, similar to the static sound of an untuned radio. While its impact can be reduced by averaging multiple measurements, it cannot be eliminated entirely. In electronic circuits, white noise typically has thermal origin (Johnson noise), or is due to the discrete nature of electrical charges (shot noise).

2. $1/f$ **(Pink) Noise:** This noise density increases as the frequency decreases, meaning it is more dominant in slow, low-frequency measurements. It is particularly troublesome for precision DC measurements. There is no single origin of $1/f$ noise - it has been observed in a very wide range of physical systems, including resistors and semiconductors.

3. **Random walk** ($1/f^2$)**:** This represents very slow, long-term fluctuations in the measurement reading, often caused by temperature and humidity changes or component ageing, primarily from voltage reference instability and offset instability. Random walk appears as drift-like behaviour in time series.

Integrating ADCs are designed to reduce noise by averaging the signal over a specific period (integration time). However, there is a trade-off: longer integration times reduce noise but slow down the measurement speed (aperture time). For example, the industry-standard 3458A type Digital Multimeter (DMM) reaches a performance limit where $1/f$ noise begins to dominate at integration times between 100 ms and 1 s [1]. The newer 8588A type DMM has successfully mitigated this low-frequency $1/f$ noise, but at the cost of increasing the white noise level (higher noise spectral density) [2].

3.2 Linearity

Ideally, a digitiser's output should be perfectly proportional to its input. If the input voltage doubles, the digital reading should exactly double. Any deviation from this perfect straight-line relationship is called non-linearity.

Linearity is compromised by imperfections in electronic components (resistors, capacitors, operational amplifiers) and the conversion process itself. These components may behave differently under varying voltages or temperatures (parameter drift).

Current state-of-the-art performance benchmarks indicate:

- Integrating ADCs (IADC): Excellent linearity of approximately 0.05 ppm (parts per million) for DC measurements [3].
- SAR ADCs: Lower linearity of approximately 0.9 ppm [4].

It is important to note that linearity specifications can vary based on how they are defined (e.g., relative to full scale or ± full scale), which can lead to a factor of two difference in reported values. In this report, we report linearity relative to full scale.

3.3 Timing Accuracy

Timing accuracy refers to the digitiser's ability to take a measurement at the exact moment requested by an external trigger signal but also to keep accurate timing interval between successive sampling events when they are dictated either by external trigger or by internal timing.

Integrating ADCs (IADCs) operate on a complex, continuous internal timing cycle that cannot be easily interrupted. Because the internal cycle must run smoothly to maintain accuracy, the digitiser cannot always start a measurement instantly when an external trigger is received. This results in a variable delay, known as jitter.

For example:

- The 3458A DMM has a jitter of approximately 6 ns [5].
- The 8588A DMM has a significantly larger jitter of 100 ns [2].

This timing uncertainty limits the ability to synchronize measurements precisely with external events or other instruments. The jitter also increases the white noise when a large-scale sine wave is sampled.

4 Mitigations

To address these barriers, several mitigation strategies were studied and are proposed in this report.

4.1 Parallel Operation

White and $1/f$ noise can be reduced by operating multiple digitiser stages in parallel. Ideally, this reduces uncorrelated noise by a factor of $\sqrt{N}$ where N is the number of digitisers. In certain conditions, harmonic distortions can also be reduced. An effort to significantly reduce these distortions has shown mixed results, demonstrating that this is a very challenging (or perhaps impossible) problem to improve significantly.

4.1.1 Three parallel SAR digitisers in averaging mode

Experiments with SAR ADCs (Fluke 8588A) in averaging mode have demonstrated SNR improvements consistent with this theory (approx. 4 dB improvement with 3 units) [2]. Each digitiser was employing averaging technique over time (t_{ad} which corresponds to the sampling time used multiplied by a number of samples used in an average) to mimic the IADC sinc roll-off frequency response due to integration of input signal over the aperture time (t_a, which is the integrating time for the input signal).

Parallel operation demonstrated not only reduced white noise, but also reduced harmonic distortion of the combined result from three digitisers operating in parallel, as shown in Table 1. However, the harmonic distortion reduction was demonstrated only at the 102.6 Hz and 1 kHz input signal frequency, but not at the 19.8 kHz input frequency.

Table 1: SNR results using APx555 analogue source. The SAR ADC was operating at 1 MHz sampling frequency, averaging each result using five samples. *Results given for SAR ADC operating at its full speed 5 MHz sampling frequency.

Fluke 8588A	102.6 Hz		1 kHz		19.8 kHz		19.8 kHz*	
	SNR dB	THD dB	SNR dB	THD dB	SNR dB	THD dB	SNR dB	THD dB
DMM1	80.15	-104.72	80.13	-109.84	80.10	-102.57	79.21	-102.75
DMM2	80.23	-104.47	80.22	-109.24	80.17	-102.31	80.24	-102.18
DMM3	80.11	-104.72	80.10	-110.16	80.06	-102.63	75.65	-102.40
Average	**84.15**	**-108.12**	**84.04**	**-112.36**	**84.05**	**-102.74**	**82.22**	**-102.85**

4.1.2 Four channel SAR digitiser

Further, measurements performed with a 4-channel SAR digitiser (CERN) have also confirmed this expectation in practice.

Measurements were performed on an FGC 3.2 ANA200 board. The board, developed by CERN, consists of four independent input channels and four ADCs (LTC2378). There are five boards available at CERN and they were already thoroughly evaluated for DC performance [6]. However, the parallel operation was not studied. The signal source was a JanasCard ultra low-distortion oscillator, operating at approx. 1 kHz. The amplitude was adjustable using the analogue potentiometer.

Measurements were recorded at CERN at 1 V, 5 V and 10 V amplitudes. All four channels were connected in parallel. The sample rate was 500 kHz. 10 000 000 samples were acquired at each voltage, resulting in 20 s sampling bursts. As the analogue free-running 1 kHz oscillator was used as a signal source, synchronisation was not possible and all further analysis was performed taking into account non-coherent asynchronous sampling.

All results were calculated as an average from ten results, obtained from each data stream divided into ten equal length parts. Therefore, each individual result was derived from 1 000 000 samples.

Signal to noise ratio (SNR) SNR was calculated as a ratio between signal (1 kHz) voltage and noise voltage, which was derived from median linear spectral density value, multiplied by the square root of spectrum frequency band [7]. Results are given in Table 2. Here, it is clear that the average SNR is improved by a factor of $\sqrt{4}$ or 6 dB.

Voltage V	Ch1 dB	Ch2 dB	Ch3 dB	Ch4 dB	Avg dB
10	96.65	96.69	96.69	96.71	102.55
5	90.24	90.19	90.22	90.20	96.19
1	76.95	76.90	76.81	76.97	82.93

Table 2: Signal to noise ratio (SNR) measured at three different input voltages.

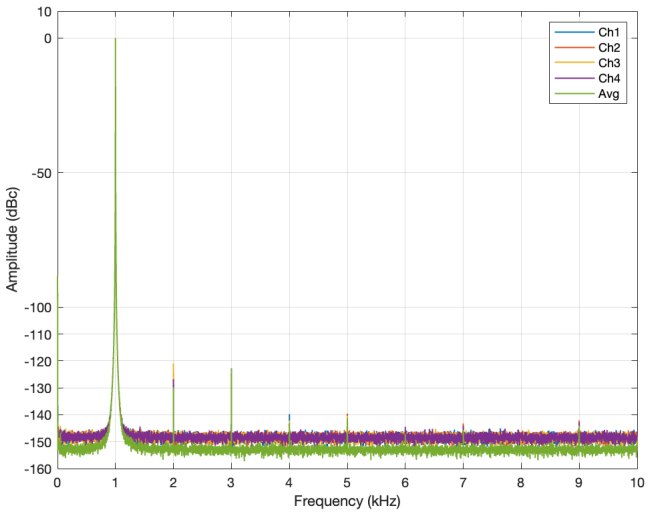

Figure 1: Spectrum of all four channels and the average. $V = 10\,\mathrm{V}$. A Hanning windowed FFT was used.

Harmonic distortion All harmonic distortions are calculated using a multi-harmonic sine fit procedure [8]. Results are given in Tables 3 to 5.

h	1	2	3	4	5	6	7	8	9	10
Ch1	0.0	-126.6	-122.7	-140.6	-141.2	-148.3	-144.6	-155.1	-142.5	-149.9
Ch2	0.0	-127.3	-123.0	-143.1	-139.8	-148.0	-143.3	-153.2	-141.7	-153.9
Ch3	0.0	-120.7	-123.0	-144.3	-140.5	-145.8	-145.4	-157.8	-143.5	-151.0
Ch4	0.0	-127.0	-123.6	-145.9	-142.7	-147.2	-144.9	-156.3	-142.9	-152.3
Avg	0.0	-129.7	-123.1	-143.5	-141.1	-147.4	-144.6	-156.4	-142.7	-151.8

Table 3: Harmonic components expressed in dBc, for the fundamental tone and first nine harmonics for input voltage $V = 10\,\mathrm{V}$

.

Harmonic phases In Table 3, it can be seen that the second harmonic is slightly decreased for the averaged data. The phase information given in Fig. 2, 3 and 4 explains this finding by showing that

h	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Ch1	0.0	-128.1	-127.9	-161.2	-136.1	-157.3	-137.8	-152.0	-144.2	-149.2
Ch2	0.0	-130.9	-126.7	-157.7	-132.2	-158.1	-136.1	-151.5	-141.6	-156.3
Ch3	0.0	-125.1	-128.2	-152.2	-135.4	-159.8	-135.8	-158.0	-144.4	-152.0
Ch4	0.0	-130.1	-129.9	-158.5	-138.6	-151.8	-139.2	-155.8	-144.6	-159.2
Avg	-0.0	-130.4	-128.1	-163.2	-135.3	-166.3	-137.1	-164.4	-143.7	-159.6

Table 4: Harmonic components expressed in dBc, for the fundamental tone and first nine harmonics for input voltage $V = 5\,\mathrm{V}$

.

h	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Ch1	0.0	-137.4	-130.2	-154.8	-141.2	-142.9	-142.4	-144.1	-141.3	-149.0
Ch2	0.0	-136.3	-130.0	-146.2	-136.6	-148.9	-137.7	-145.4	-143.0	-148.2
Ch3	0.0	-135.2	-129.6	-139.8	-139.8	-139.7	-139.0	-140.2	-145.3	-148.0
Ch4	0.0	-150.4	-134.0	-143.1	-136.5	-146.7	-139.7	-141.6	-145.2	-157.0
Avg	-0.0	-141.0	-130.8	-149.1	-138.6	-148.6	-140.9	-156.7	-143.8	-166.3

Table 5: Harmonic components expressed in dBc, for the fundamental tone and first nine harmonics for input voltage $V = 1\,\mathrm{V}$

.

second harmonic distortion phase can deviate significantly among ADC channels, which is not the case for other harmonics.

The signal phase analysis revealed that by using inverted approach, where two channels would be connected directly and two channels in reverse connection, the even harmonics would be largely cancelled at all voltage levels, as the channel harmonic phases are very similar and their amplitudes are also similar (see App. A). These measurements could not be conducted reliably due to interference effects in the measurement setup. However, this could provide a fruitful area for further investigation.

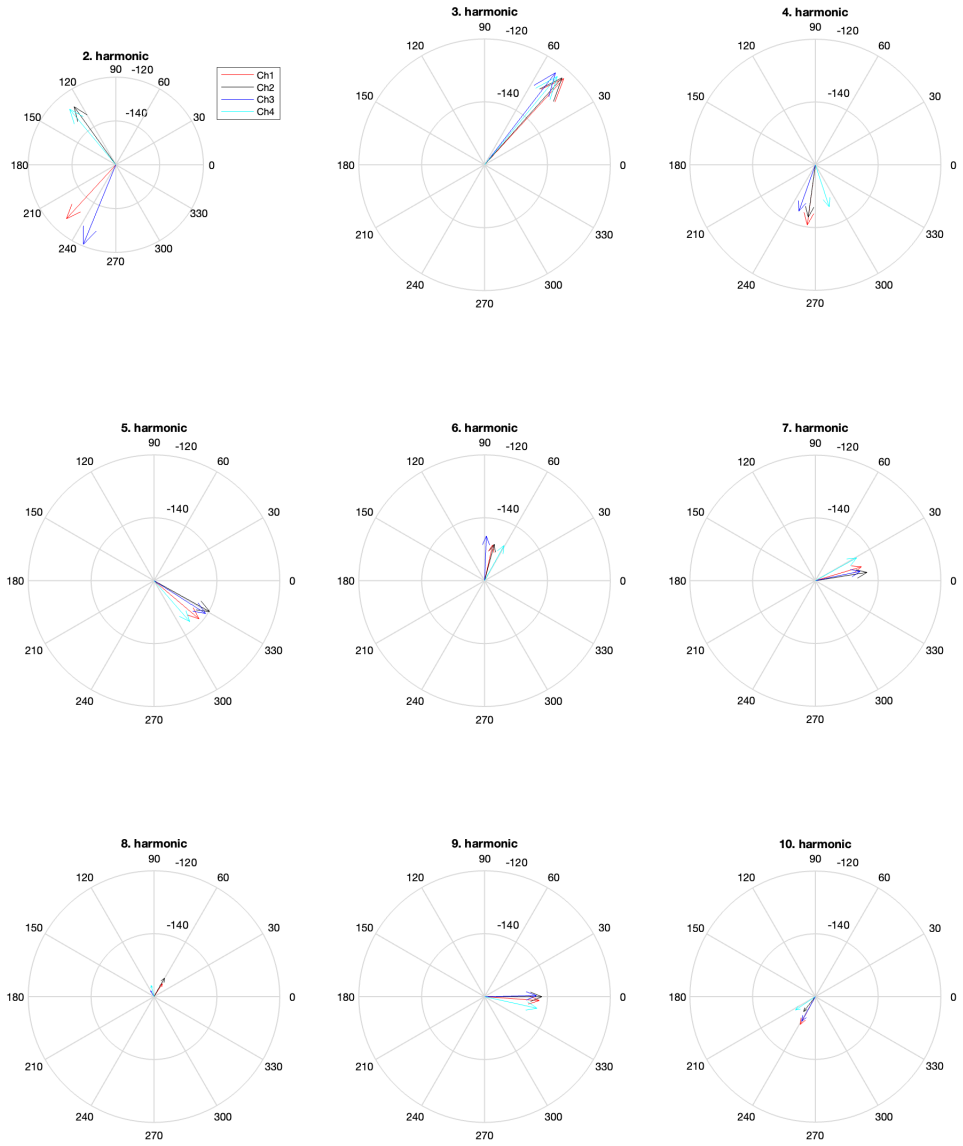

Figure 2: Harmonic phases at 10 V input amplitude, relative to fundamental signal.

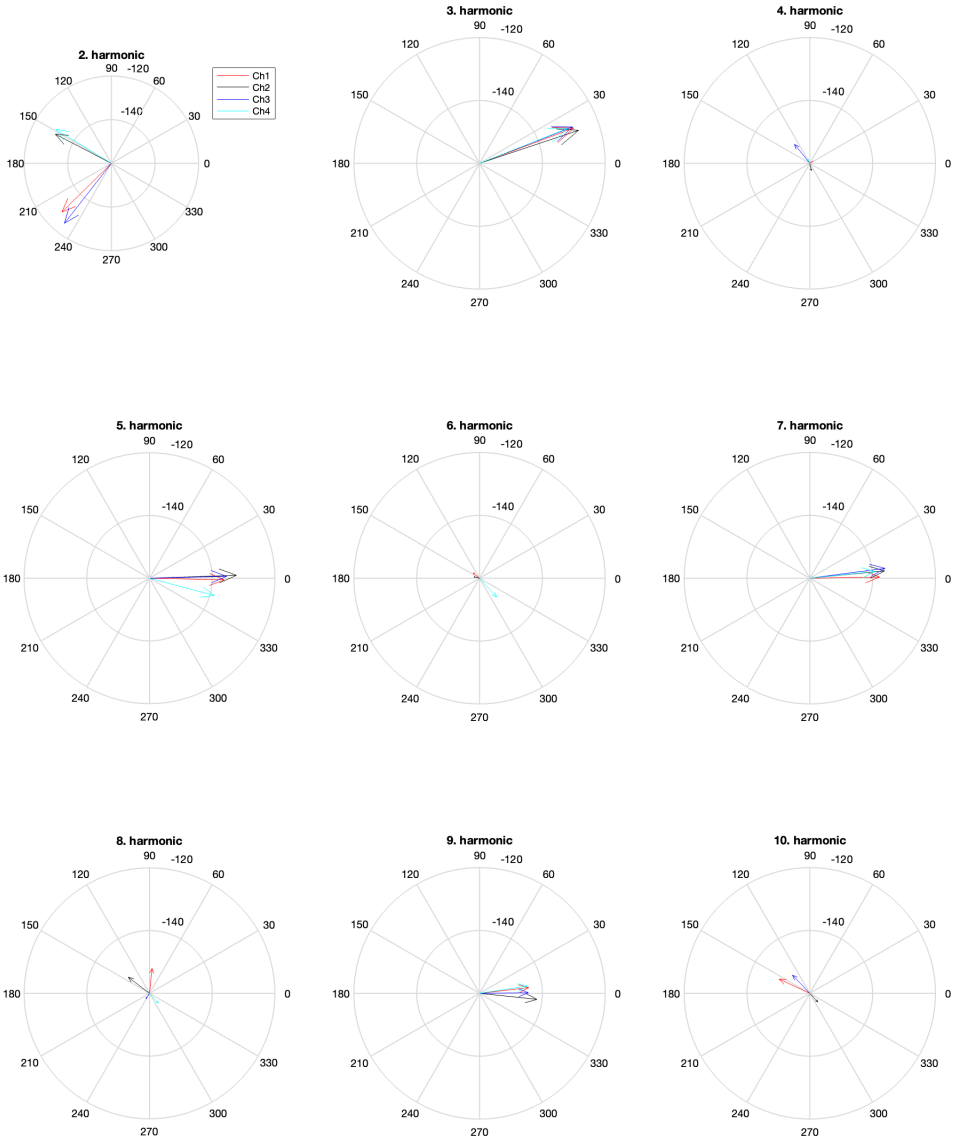

Figure 3: Harmonic phases at 5 V input amplitude, relative to fundamental signal.

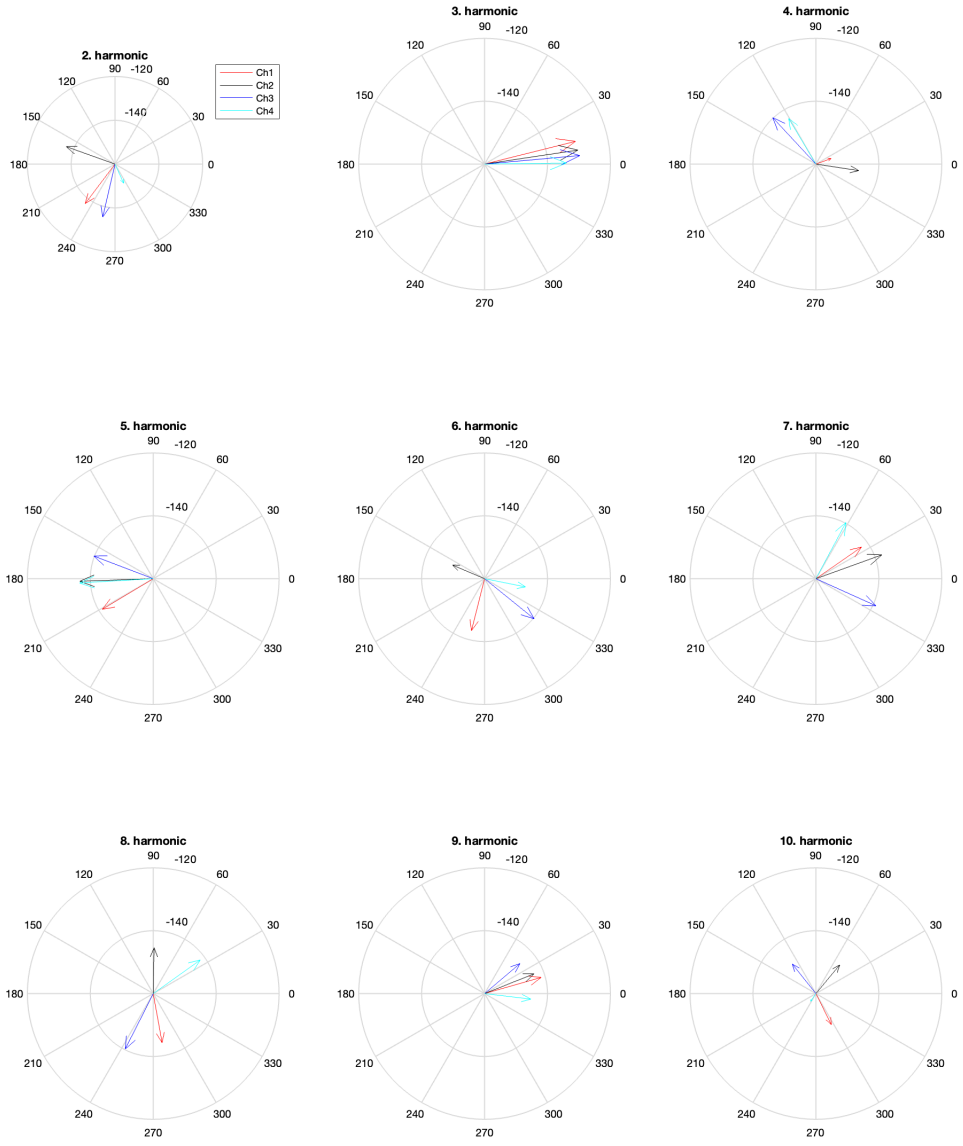

Figure 4: Harmonic phases at 1 V input amplitude, relative to fundamental signal. Scattered phases at certain harmonics are due to a strong influence of noise in these harmonic measurements.

This was the second attempt to characterise parallel SAR ADC operation. Similar to the results reported in Section 4.1.1 the results of this test confirm that the signal to noise ratio is improved by the expected amount of $\sqrt{N}$ where N is the number of channels.

Further analysis has revealed that parallel channels using the same components exhibit quite some variation in their harmonic distortion performance. This variation can be seen in the different harmonic component levels. The second harmonic phases were found to be inconsistent among four channels, while this is not the case for higher harmonics. From this observation we can conclude that, for simple averaging, the second harmonic can be somewhat compensated using parallel operation, while the higher harmonics will remain at the average level among all channels.

4.1.3 DC Offset Separated Averaging

This study focuses on the use of successive approximation register (SAR) ADCs configured in parallel, aiming to enhance the performance of precision digitisers by exploiting averaging and cancellation techniques at the hardware level.

This concept leverages the statistical averaging of uncorrelated noise sources from multiple ADCs, as well as the cancellation of deterministic non-linearities, particularly even-order harmonics, through signal inversion strategies. Simulation work includes the testing of configurations where the ADC inputs are offset in DC to break the periodic nature of integral nonlinearity (INL), and where signal inversion and averaging pairs are used to suppress even harmonic distortion. These simulations serve to establish the design parameters for an experimental prototype, and to assess the feasibility and potential benefits of such approach.

Parallel operation setup Here we present a simulation-based investigation of a digitiser architecture employing parallel SAR ADC channels, organised in two pairs. One pair receives the input signal directly, while the other pair is fed the inverted signal. Each of the ADCs in a channel is further distinguished by a unique DC offset referred to its input. The goal of this approach is to break the spatial periodicity of ADC integral nonlinearity (INL) and, through strategic averaging, to suppress both random and deterministic error components. Clearly, this approach comes at the cost of reducing the effective usable input range, which is reduced by the largest DC offset applied.

The block diagram in Fig. 5 outlines the simulation architecture, which serves as the conceptual basis for the following analysis.

The motivation stems from the need to reduce INL-induced harmonic distortion and to increase signal-to-noise ratio (SNR) by leveraging the statistical independence of the offset-modulated digitisation paths.

Theory Behind Averaging and Inverted Operation Averaging the outputs of multiple ADCs with independent INL and noise sources improves both noise performance and linearity. When M ADCs are averaged together, the quantisation noise and uncorrelated random noise are reduced by a factor of $\sqrt{M}$, assuming ideal averaging.

Applying different DC offsets to parallel ADCs breaks the periodic alignment of their INL characteristics, which otherwise tend to introduce spurious tones at predictable locations. The combined output thus averages out deterministic non-linearities to a significant degree. The DC offsets described here are determined using

$$\mathrm{DC}_i = -V_{\mathrm{DC}} + \frac{2V_{\mathrm{DC}}}{m-1} \cdot (i-1), \quad \text{for } i = 1, 2, \ldots, m, \tag{1}$$

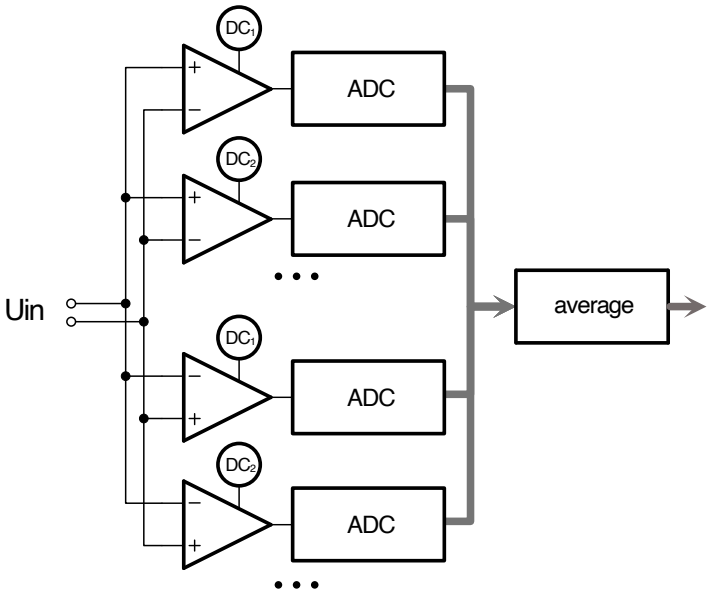

Figure 5: Block diagram of the simulated parallel digitiser architecture with DC offset separation and signal inversion. The pair structure can be expanded further for both direct and inverted connection.

where m is the number of ADC paths used for either direct or inverted connection and V_{DC} is the chosen maximum DC offset. The available signal input range is reduced by the value of V_{DC}.

Furthermore, signal inversion in half of the channels suppresses even-order non-linearities. Let $x(t)$ be the input, and $f(x)$ be the non-linear response of the ADC. Then the average of $f(x)$ and $f(-x)$ eliminates the even-order terms in the Taylor expansion of $f(x)$:

$$\frac{f(x) - f(-x)}{2} = \text{odd-order terms only.} \tag{2}$$

This principle enables suppression of the second, fourth, and other even harmonics, which are often dominant in real SAR ADCs due to capacitor mismatch and comparator asymmetries.

The combination of DC offset separation and signal inversion enables spatial dithering of nonlinearity error terms, yielding a more linear and noise-suppressed effective digitisation chain. In the following sections, we demonstrate this through simulation results using available measured and published INL datasets. Of course, this works as explained only if the non-linearities of all ADCs and their input driving stages are exactly equal.

Digitisation Function Overview All simulations presented in this chapter are based on a general-purpose MATLAB function `digitize`, developed to model the behavior of SAR ADCs with high configurability. This function supports multi-channel simulations and provides comprehensive control over the ADC parameters.

Some key capabilities of the `digitize` function include:

- Support for multiple simultaneous ADC channels through a user-defined function handle returning a matrix of input signals.

- Specification of ADC resolution for each channel independently (bit depth).
- Definition of the reference voltage or full input range.
- Sampling rate selection via the sampling time `Ts` or number of samples `N`.
- Support for sample-to-sample jitter to model clock timing variations.
- Injection of white and $1/f$ noise through configurable noise spectral density and flicker parameters.
- Optional linear drift modelling per channel.
- Application of integral nonlinearity (INL) either from a lookup table or from a user-defined polynomial (via the `polynom` field).
- Simulation of differential nonlinearity (DNL) using table-based or interpolated models.
- Selection of interpolation method (linear or spline) for INL/DNL application.
- Per-channel DC offset assignment to enable offset-separated operation.
- Grouping of multiple ADC instances per logical channel to model averaging of parallel converters.

This flexibility enables detailed exploration of how INL, offset, noise, and averaging affect the digitised output. All simulations shown in the subsequent sections rely on this framework, using both measured and synthetic INL datasets for validation.

```
% DIGITIZE Simulates an ADC with nonlinearities, noise, and imperfections.
% All input parameters are passed via a structure.
% Output includes quantized codes, reconstructed voltages, time vector,
% and optionally interpolated INL/DNL tables.
%
% function output = digitize(input)
%
% INPUT structure fields:
% .func         : @(t) function handle returning [n_channels x N] array of input signals
% .Ts           : Sampling time [s]
% .n_bits       : Scalar or vector [n_channels] specifying ADC resolution [bits]
% .t_end or .N  : Total simulation time [s] or number of samples
% .Vref         : ADC reference voltage [V] (optional if .fullscale is provided)
% .fullscale    : Full input voltage range [V] (optional if .Vref is provided)
%
% Optional fields:
% .offset       : Scalar/vector/cell, DC offset per channel [V]
% .polynom      : Cell array of polynomial distortion coefficients per channel
% .nsd          : Scalar/vector/cell, white noise spectral density [V/sqrt(Hz)]
% .nf1          : Scalar/vector/cell, 1/f noise coefficient
% .nf1_bw       : Bandwidth cutoff for 1/f noise [Hz]
% .drift        : Scalar/vector/cell, drift rate [V/s]
% .jitter       : Sampling jitter standard deviation [s]
% .inl          : Vector or cell array, INL tables per channel
% .dnl          : Vector or cell array, DNL tables per channel
% .inl_unit     : 'lsb' or 'ppm' (default 'lsb')
```

```
% .dnl_unit     : 'lsb' or 'ppm' (default 'lsb')
% .inl_interp   : 'linear' or 'spline' (default 'linear') for INL interpolation
% .debug        : true/false, enable debug mode (default false)
%
% OUTPUT structure fields:
% .digital      : [n_channels x N] array of quantized ADC codes
% .voltage      : [n_channels x N] array of output voltages
% .t            : Time vector [1 x N]
% .INL          : [n_channels x N] array of applied INL errors [V]
% .DNL          : [n_channels x N] array of applied DNL errors [V] (only if DNL
    provided)
```

Listing 1: Header help section of the `digitize` function.

4.2 Results for AD4630 INL

The first set of simulation results focuses on the AD4630-24 ADC, using a representative INL table extracted from the manufacturer datasheet (Fig. 6). The INL table was extended across all code values using spline interpolation to ensure continuity and resolution beyond the original measurement points.

Four ADCs were simulated in parallel. Two channels received the signal directly and two channels received the inverted version of the same signal. Within each pair, the two ADCs were assigned different static DC offsets to intentionally break the INL periodicity. In this case, $DC_1 = -DC_2$. The DC offsets were varied between 0 and 0.6 times the full input range (5 V) to explore the impact on linearity. The sampling frequency was set to 500 kHz, and the input signal was a 1 kHz sine wave. A total of 100 periods were sampled, and the bit depth was set to 32 bits to eliminate quantisation noise effects and isolate INL-induced distortion.

Fig. 7 shows (a) the output spectrum of a single ADC channel showing prominent INL-related harmonic distortion and (b) the spectrum of the averaged output signal after combining the four ADCs with inversion and DC offset separation.

Fig. 8 shows a comparison of residual INL amplitude as a function of DC offset magnitude for a single channel, four averaged ADC paths without inversion, and two double ADC paths using direct and inverted signal connection. Note that the DC offsets for the non-inverted configuration (four different DC offset values) are different from the inverted configuration (two different DC offset values, used in direct and inverted group).

These figures clearly demonstrate the benefit of this architecture in reducing both harmonic distortion and DC INL.

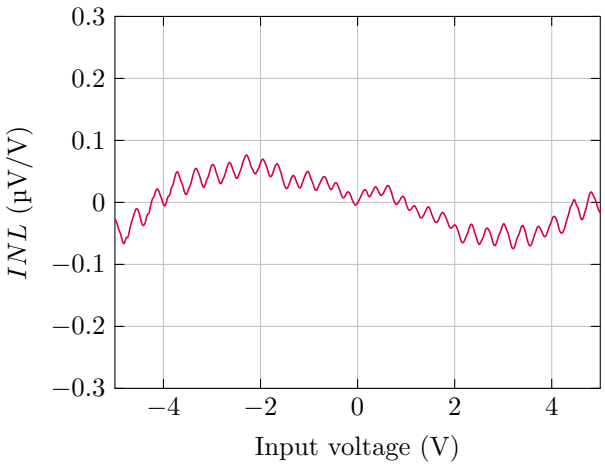

Figure 6: AD4630-24 INL plot used for this simulation.

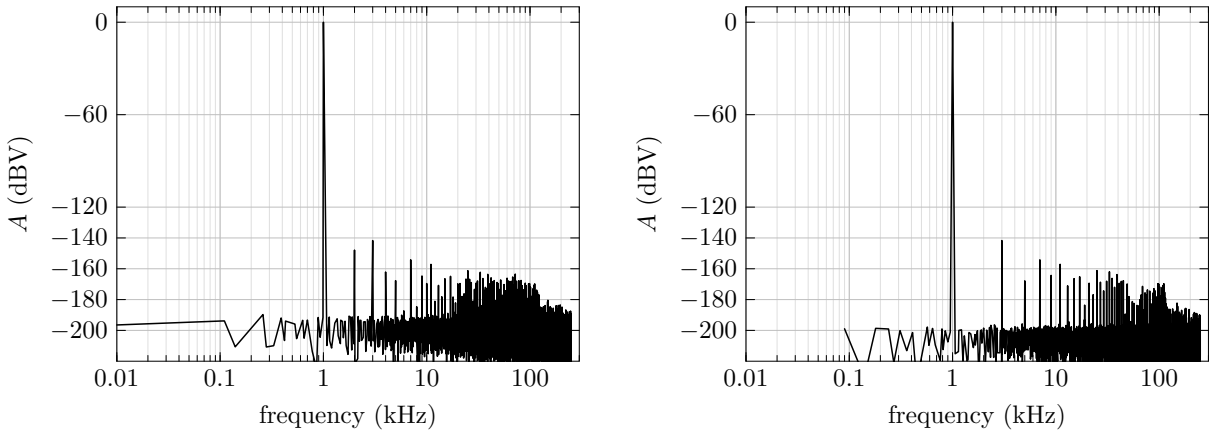

(a) Single ADC channel spectrum.

(b) Averaged output spectrum using DC offsets and inversion, using four ADC channels.

Figure 7: Comparison of output spectra for a single ADC channel and the averaged structure shown on Fig. 5. Removal of even-order harmonics is demonstrated due to the inverted channels configuration.

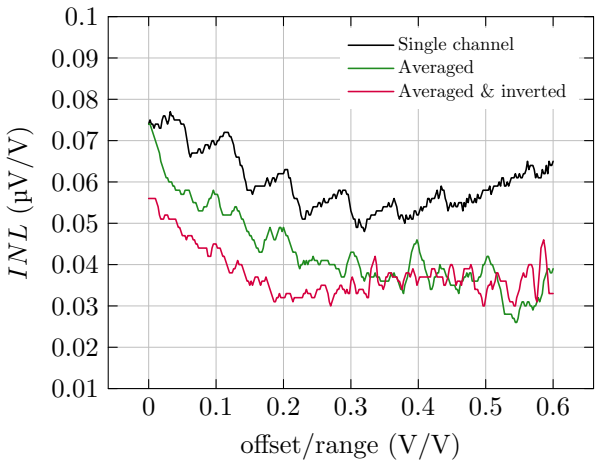

Figure 8: Integral nonlinearity (INL) for a single channel and the averaged output structure. Averaging with DC offset separation and inversion reduces peak INL amplitude.

Custom IADC INL The second set of simulation results focused on the custom IADC, developed by the collaborator Jaromir Sukuba, using a representative INL table provided by the developer (Fig. 9). The INL table was extended across all code values using spline interpolation to ensure continuity and resolution beyond the original measurement points. The only difference to the AD4630 simulation was the use of 10 V input range.

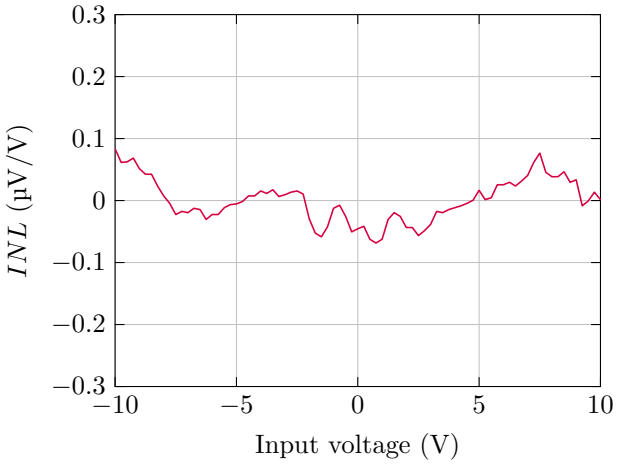

Figure 9: Jaromir IADC INL plot used for this simulation.

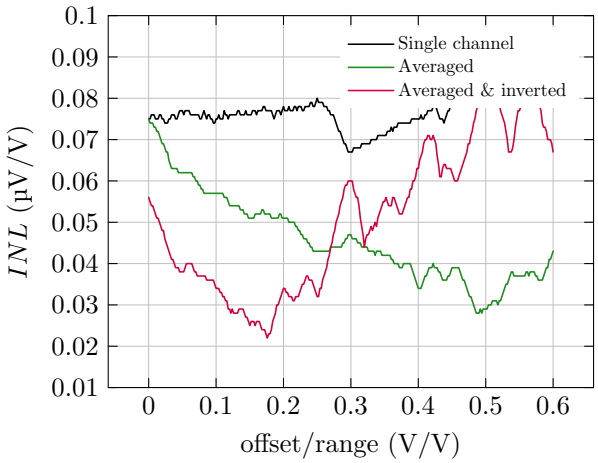

Figure 10: Integral nonlinearity (INL) for a single channel and the averaged output structure. Averaging with DC offset separation and inversion reduces peak INL amplitude.

Fig. 10 plots the resulting INL. Single channel shows the relative nonlinearity result for a single channel as a function of the reduced input span. Averaged result shows four identical IADC in parallel with offset set to (-offset, -offset/3, offset/3, offset) V, where offset is varying from 0 to 6 V. Accordingly, the sine amplitude is reduced to prevent clipping any of the IADC under test.

Conclusion This study has demonstrated that employing an averaged configuration composed of both direct and inverted ADC groups can significantly improve integral nonlinearity (INL), particularly when small to moderate DC offsets are applied. Compared to a configuration using the same number of directly connected ADC paths with twice as many unique DC offsets, the mixed (inverted + direct) architecture shows a more rapid reduction in INL for small offsets.

However, as the DC offset increases — approaching approximately 25 % of the ADC's input range — the benefit of signal inversion diminishes. At these larger offsets, both approaches offer similar improvements, and in some cases, the simpler direct-only averaging may outperform the inverted configuration.

It is important to note that these findings are based on simulations assuming ideal conditions, where all ADC paths are identical apart from their additive noise. In practical applications, where real ADC channels exhibit individual INL characteristics, the effectiveness of either method may vary.

Finally, the optimum set of DC offsets for INL suppression cannot be determined analytically in the general case. Instead, it must be established through empirical measurements of the actual converter behaviour.

The next option is to test inverting connection, where two channel are connected to an inverted signal (this is possible due to balanced input configuration) and the combination signal is obtained by a difference. This way, as it is generally known, the even order harmonics can be significantly suppressed (see App. A). Further, a DC offset can be applied to each channel (say -15 %, -5 %, +5 % and +15 % or range for each channel), which will break periodicity of the ADC distortion and potentially improve linearity when averaged.

4.3 Input Reversal (Chopping):

$1/f$ noise and random walk can be mitigated using input reversal techniques. By measuring the signal in positive and reversed polarity, the average effectively removes systematic bias (random walk and pink noise) originating in the input circuits and ADC [9]. Implementation options include using auto-zero based input stages.

$1/f$ noise and random walk can both be effectively removed using an input reversal technique. The measured signal is connected to the input in one polarity for the first measurement and in reversed polarity for the second measurement. The average of both results then effectively removes $1/f$ noise and/or random walk, which both represent systematic bias from the reference zero value. White noise does not provide any bias and is therefore not removed using this technique. Another way to describe the principle is to state that white noise has no temporal auto-correlation, while other mechanisms such as $1/f$ noise have non-zero correlations.

Possible configurations for IADC are shown on Figure 11.

- Both variants **A** and **B** need a precision inverter, which probably has to be realized using a similar (or the same) COPA as the input stage.
- Variant **B** is more economical, but it breaks the concept of the current-steering switches – the polarity reversal should be done with a switch that tolerates high voltages (unlike the current steering switch elements (SW_{in}) that do not operate at higher voltage).
- Variant **A** keeps the current-steering topology, but it adds two more elements that need matching (two switches and the two resistors). The matching could be checked periodically by coupling both signals simultaneously and measuring their difference. It could be accomplished either using the main integrator, or using a parallel path that diverts the current instead of coupling it to GND.

A. Single-ended input + inverter + 2 resistors + 2 current-steering switches

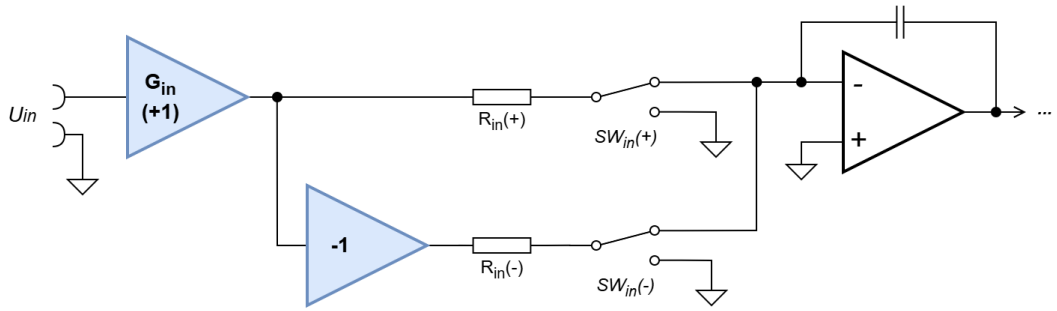

B. Single-ended input + inverter + voltage switch

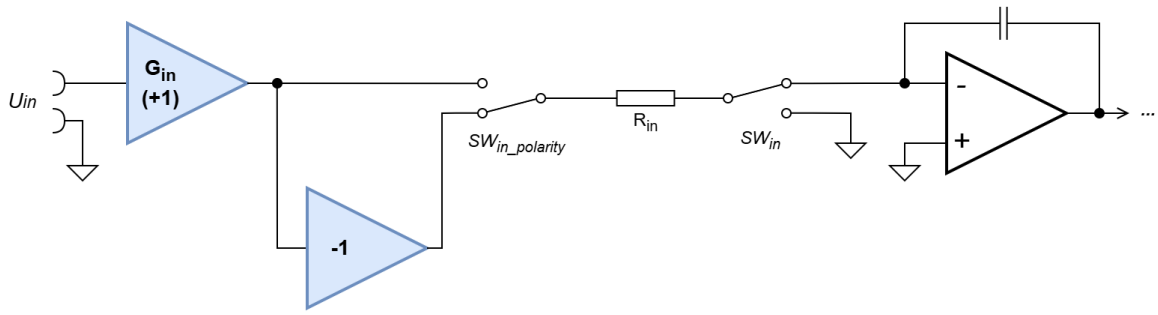

C. Differential input + voltage-reversing switches

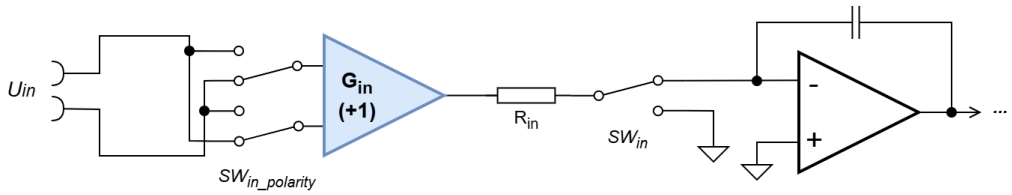

Figure 11: Three schemes for input switching implementation.

- Variant **C** needs a stage with differential inputs and single-ended output. The advantage is that it relaxes the requirements on the isolation of the PSU, and it can be used with balanced line (e.g. shielded twisted pair), so it would be more immune to electromagnetic interference in real operating conditions. The disadvantage is that the input signal source will see charge injection from the polarity reversal switches. One mitigating option is to implement the switches after a buffering stage, in order to isolate them from the source. The disadvantage is the need to use two precision buffers, in addition to the precision difference amplifier.
- It should be noted that in any case, good buffering and inversion of U_{ref} is needed. The dynamic performance of these stages is not important, they only need to have good DC stability.

4.4 Timing Correction:

A Time Interval Counter (TIC) can be added to measure the delay between the low-jitter internal timing and the high-jitter external trigger (Fig. 12). Using this data, the integrator voltage at the exact external trigger event can be reconstructed via interpolation. The main innovation proposed is adding a time

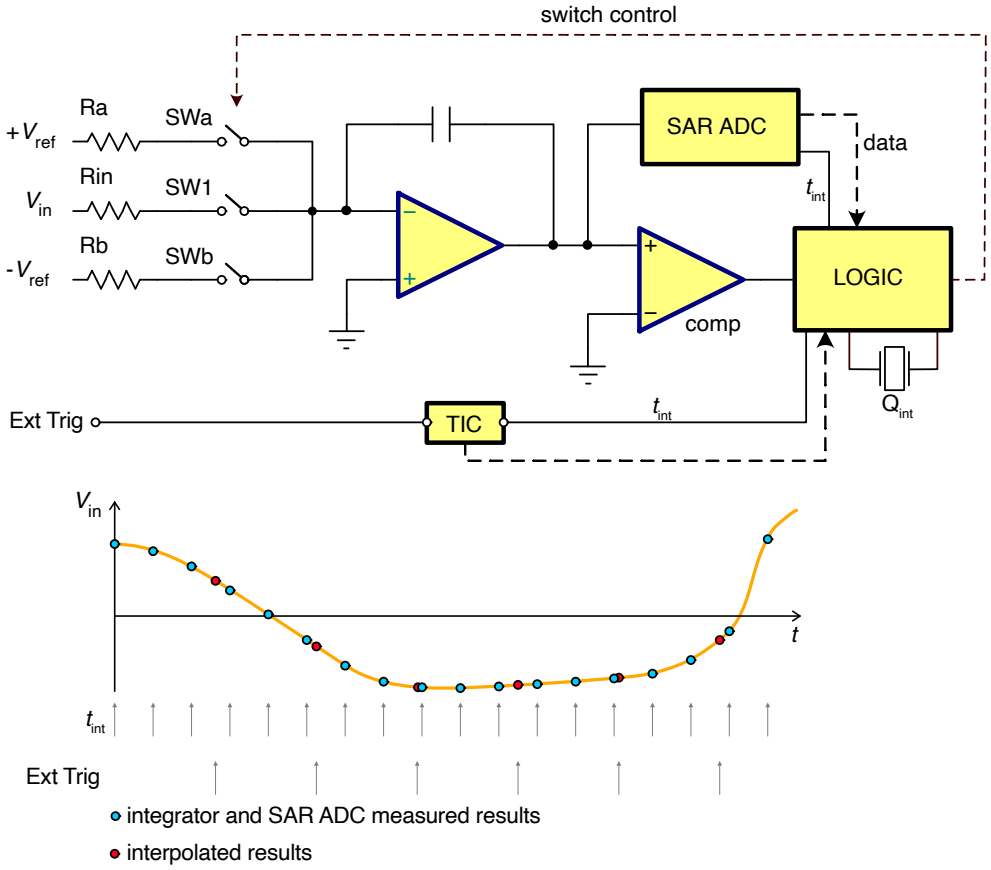

Figure 12: Proposed concept for external trigger timing solution.

interval counter (TIC) to measure timing relationship between internal timing (low jitter) and external trigger signal (high jitter). With adequate signal processing and an accurate TIC, this scheme can provide for accurate external trigger and can further improve on a jittered external trigger. The IADC operates in a manner to constantly seek balance of the integrator output voltage. Logic controls the input switches during the run-up with a multi-slope pattern that keeps the number of switch transitions constant [3]. A single pattern defines an internal cycle. SAR ADC reads the integrator output voltage at each completed internal cycle and the result is stored in internal memory. With this, the starting value (and the end value) of each cycle are preserved, thus the residual information (apart from counting the number of reference currents during the tint intervals) is fully preserved at each cycle. The integrator capacitor reset switch is therefore not needed, and the rundown procedure is also not needed.

Without timing corrections (interpolation), the following operations are possible:

- Continuous run (internal timer)
- During the time results are stored, it is possible to restore any aperture time information, down to one tint interval. Therefore, the measurement need not be repeated to obtain result at different aperture time.
- External trigger has a maximum jitter of $\pm t_{\text{int}}/2$ (rectangular probability).

With timing corrections (interpolation), the following improvements are possible:

- TIC measures a delay between external trigger pulse and internal timing cycle. Using interpolation, the delayed point can be retrieved from the internal SAR ADC and integrator switching pattern data using interpolation. This can significantly reduce trigger jitter (theoretically down to TIC jitter and internal timing jitter, combined).
- Further, for continuous external triggering for each sample, the external trigger signal jitter can be reduced by postprocessing and calculating "ideal" trigger timing delays from average external trigger timings. However, starting trigger timing cannot be improved and is taken as a single event.

However, the implementation of TIC is highly challenging and not cost effective. A similar solution, but with different processing, could be realised with a hybrid digitiser shown on Fig. 32, where the Wideband SAR ADC provides accurate trigger points and IADC trigger is delayed by a highly jittered delay. In this topology, the IADC might run entirely on the internal timer and the information comparison between SAR ADC results and IADC results can provide accurate time adjustment information.

4.5 Component Selection

Linearity can be improved by selecting critical passive components (resistors, capacitors) of the highest quality, understanding switch imperfections and composite operational amplifier (COPA) [10] influence on the IADC based digitiser design. This selection shall be supported by thorough analysis, that shows component particular imperfection on linearity influence. The model described in Section 5 can be used to investigate the effect of the imperfections of components on the performance of the IADC and consequently guide the IADC developer in the selection of components. The component analysis is reported in [10] and resistor non-linearity measurement results are further reported in [11].

5 Modelling of IADC

To improve the accuracy of modern precision digitisers, specifically Integrating Analog-to-Digital Converters (IADCs), it is essential to understand the behaviour of their internal components before building the actual hardware. "Modelling" refers to the creation of a virtual, mathematical representation of the digitiser's circuitry—including its resistors, capacitors, switches, and amplifiers.

By simulating how these components interact and how their individual imperfections (such as temperature drift or non-linearity) affect the final measurement, we can predict errors that would otherwise be difficult to detect in a physical prototype. This section details the development of these models using standard tools like LTspice for circuit-level behaviour and MATLAB/Simulink for system-level performance. The goal is to identify which components are critical to accuracy and how mitigations can be optimised to suppress errors and enhance overall accuracy.

5.1 Circuit modeling using LTspice

5.1.1 Advanced COPA based integrator analysis

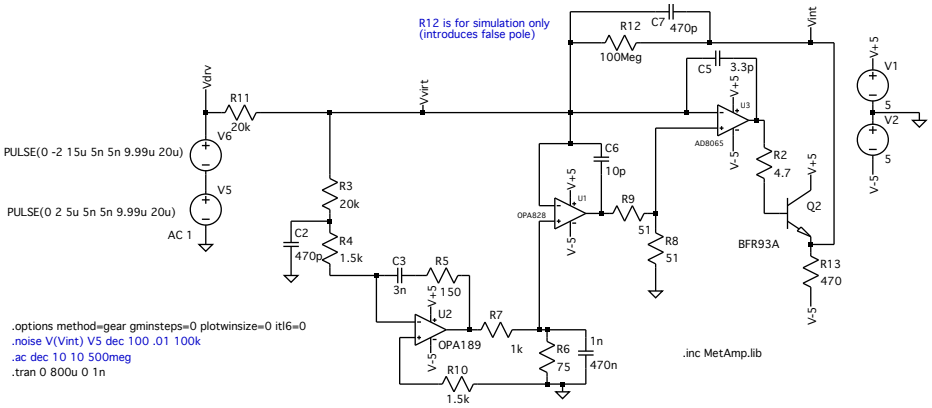

Figure 13: LTSpice simulation of new COPA as an integrator.

The long-term accuracy of the integrator is ultimately limited by unintended charge coupled into the main integrating capacitor C_7 during switching events (see Fig. 13, which gives an LTspice circuit reported in [12]). Here, we quantify this effect using transient data exported from LTspice simulations, focusing on the net residual charge per switching cycle rather than on the instantaneous charge associated with individual edges.

The long-term accuracy of the integrator can be influenced by unintended charge coupled into the main integrating capacitor C_7 during switching events. This effect was investigated using transient data exported from LTspice simulations [12] (see Fig. 13, with emphasis on the net residual charge per switching cycle, which is the quantity relevant for long-term accumulation.

Waveform observation Figure 14 shows the virtual ground voltage V_{virt} during a single switching event and 4 V input voltage change. Although the summing node is nominally held at virtual ground, a short disturbance is visible during the switching transition. This disturbance coincides with a transient

current flowing into the integrating capacitor, shown in Fig. 15. The current waveform consists of a fast spike with a duration on the order of tens of nanoseconds, followed by a weaker settling tail. The spike amplitude depends on the summing node current change, which in actual IADC means on the reference voltage and resistor defining reference current. As those values are stable throughout the conversion, the spike amplitude is also stable.

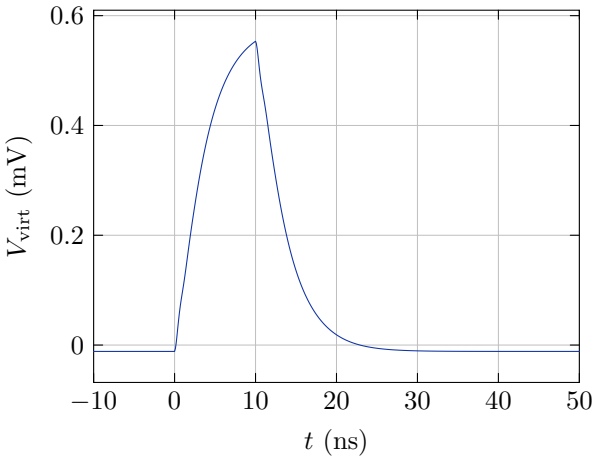

Figure 14: Virtual ground voltage V_{virt} during a single switching event, showing a short disturbance coincident with the switching transition.

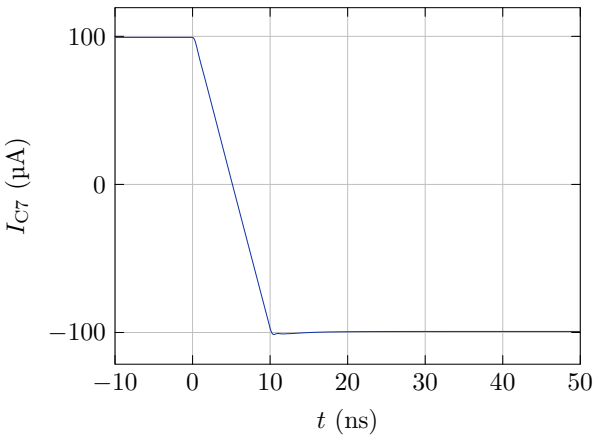

Figure 15: Current through the integrating capacitor C_7 during the same switching event, illustrating the fast spike and subsequent settling tail responsible for charge injection.

These two waveforms confirm that switching activity injects a short burst of charge into the summing node and hence into C_7, even though the average node voltage is well controlled.

Derivation of spike-induced charge The charge coupled into C_7 by a switching event was derived directly from the LTspice transient data. Two complementary approaches were used and found to be

consistent:

1. *Current-based evaluation*: the instantaneous current through the integrating capacitor, $I_{C7}(t)$, was exported and time-aligned to the known switching instants. The net charge associated with a switching event was obtained by integrating the sign-corrected current over a time window surrounding the event,
$$Q = \int I_{C7}(t)\,\mathrm{d}t.$$
Because the transient includes both a fast spike and slower settling components, individual edge charges depend on the exact integration window and are therefore not uniquely defined.

2. *Cycle-based evaluation*: to obtain a robust and physically meaningful metric, consecutive switching edges were paired to form a complete PWM cycle, and the net residual charge per cycle was computed. This approach avoids sensitivity to window placement and captures only the charge imbalance that can accumulate over time.

To improve numerical robustness, the simulations were repeated with a reduced LTspice timestep of 0.1 ns, ensuring adequate temporal resolution of the fast current spikes.

Positive and negative spike-induced charge At the finest time resolution (0.1 ns timestep), the charge associated with individual switching edges becomes well resolved. For the integrating capacitor $C_7 = 470\,\mathrm{pF}$, the spike-induced charge was found to be approximately

$$Q_+ \approx +20.17\ \mathrm{fC}, \qquad Q_- \approx -21.38\ \mathrm{fC},$$

for the positive and negative switching edges, respectively. The near-equality of these values indicates a high degree of symmetry in the switching-induced charge injection, with only a small residual imbalance remaining after a full PWM cycle.

Net residual charge per cycle When consecutive edges are paired, the net residual charge per PWM cycle converges to a stable value of approximately

$$Q_{\mathrm{res,cycle}} \approx 1.2\ \mathrm{fC},$$

with negligible cycle-to-cycle variation over the simulated interval. This corresponds to an equivalent perturbation of only a few microvolts at the integrator output. While the apparent charge associated with individual edges varies with numerical resolution and analysis method, the cycle-level residual is insensitive to these choices and therefore represents the physically relevant quantity for long-term accuracy.

Implications and limitations For an integrator full-scale swing of $\pm 5\,\mathrm{V}$ and a measurement interval of 200 ms, a relative accuracy target of 1×10^{-8} corresponds to an allowable output uncertainty of approximately 50 nV, or about 24 fC of equivalent charge on C_7. The simulated residual charge per cycle is therefore small in absolute terms. Moreover, the absence of observable cycle-to-cycle randomness indicates that the residual behaves as a deterministic offset rather than as stochastic noise, and can in principle be calibrated or corrected.

It must be emphasized that these conclusions are based solely on circuit-level simulation. The magnitude and stability of switching-induced charge coupling depend on device models, parasitic capacitances, and layout-dependent effects that are not fully captured in simulation. Experimental measurements are therefore required to confirm the influence of switching transients on C_7 in real hardware.

5.1.2 A modified charge-balancing AD algorithm with a PWM DAC

A commonly used multislope charge balancing algorithm (used e.g. in HP3458A) can be described as a single-bit ADC (a comparator at the output of the integrator and feedback 1-bit DAC), a reference current switch with a short or a long pulse. The evolution of this principle uses a multi-bit ADC on the output of the integrator – for instance in HP34401 or Multislope IV ADC used in the latest Keysight multimeters. The ADC at the output of the integrator is used as a replacement for the rundown phase. Another step forward is using the ADC also during the run-up phase and multi-bit DAC in the feedback loop. This approach isn't common in the charge balancing AD converters, but it is used in delta-sigma ADCs. The main problem with multi-bit DAC is that its linearity has a direct impact on the linearity of the whole ADC, so it has to be well below 0.1 ppm in our case. The simplest implementation of a such high linearity DAC is a PWM DAC. An example of a DS ADC using this approach is described in [13].

One of the important sources of nonlinearity in charge-balancing AD converters is the dielectric absorption (DA) of the integrating capacitor. The main source of the influence of DA is the fact that the mean value of voltage on the integrating capacitor isn't zero during the run-up phase of the AD conversion. This value also varies non-linearly with the input voltage, so the goal was to design such a position of the PWM pulse that would keep zero mean value on the integrator capacitor during the cycle, and also keep zero integrator voltage at the end of the cycle. That assures pure AC voltage without DC bias. Both conditions are fulfilled by using a centralised alternating sequence, see Fig.2 C in [14]. Keeping the voltage close to zero on the integrator output at the end cycle has another advantage – the rundown phase can be very short or can be eliminated entirely, and remaining low voltage can be measured by the AD converter. [14] uses only a 4-bit DAC, so using PWM with resolution of 20 ns (50 MHz clock), and with cycle time 50 µs, maximum duration of I_{ref} pulse is about 38 µs or 1900 clock cycles, which equals 11 bits. With $I_{\mathrm{Xmax}} = 250\,\mu\mathrm{A}$ and $C_{\mathrm{int}} = 680\,\mathrm{pF}$, maximum residual voltage is 7.3 mV. This is true only for constant or slowly/constantly varying input voltage. If the input voltage changes rapidly, the voltage at the integrator output will be greater.

Simulation An enhanced version of the integrator was use for an LTspice simulation of the described conversion method. The integration capacitor was replaced by a more complex model with additional branches $C_{\mathrm{d1}}, R_{\mathrm{d1}} \cdots C_{\mathrm{d4}}, R_{\mathrm{d4}}$, modelling the effect of dielectric absorption (see Fig. 16).

The parameters of the model are based on [15] and assume a dielectric absorption of 0.02 % taken for a polystyrene capacitor type (Table II, [15]), assuming the ceramic capacitor used will be even better, but its DA at shorter charging times is not measured yet. There are several simulation parameters:

- V_{in}: input voltage 0 V to 10 V, negative input voltage isn't implemented, but it works the same way, only the function of $V_{\mathrm{ref+}}$ and $V_{\mathrm{ref-}}$ is swapped
- T_{cycle}: fixed 50 µs
- T_{on}: width of PWM pulse – here with high time resolution, real resolution is 20 ns
- T_{s}: delay between $I_{\mathrm{ref+}}$ and $I_{\mathrm{ref-}}$ pulses – typically 0.4 µs
- T_{z}: width of opposite I_{ref} pulse – typically 0.4 µs
- N: number of cycles

Influence of integrator capacitor DA While the evaluation of the effect of the DA integration capacitor on the linearity is complicated, it depends mainly on the duration of the rundown phase. Here we chose to evaluate the magnitude of the charge stored on the parasitic capacitances $C_{\mathrm{d1}} \cdots C_{\mathrm{d4}}$. This

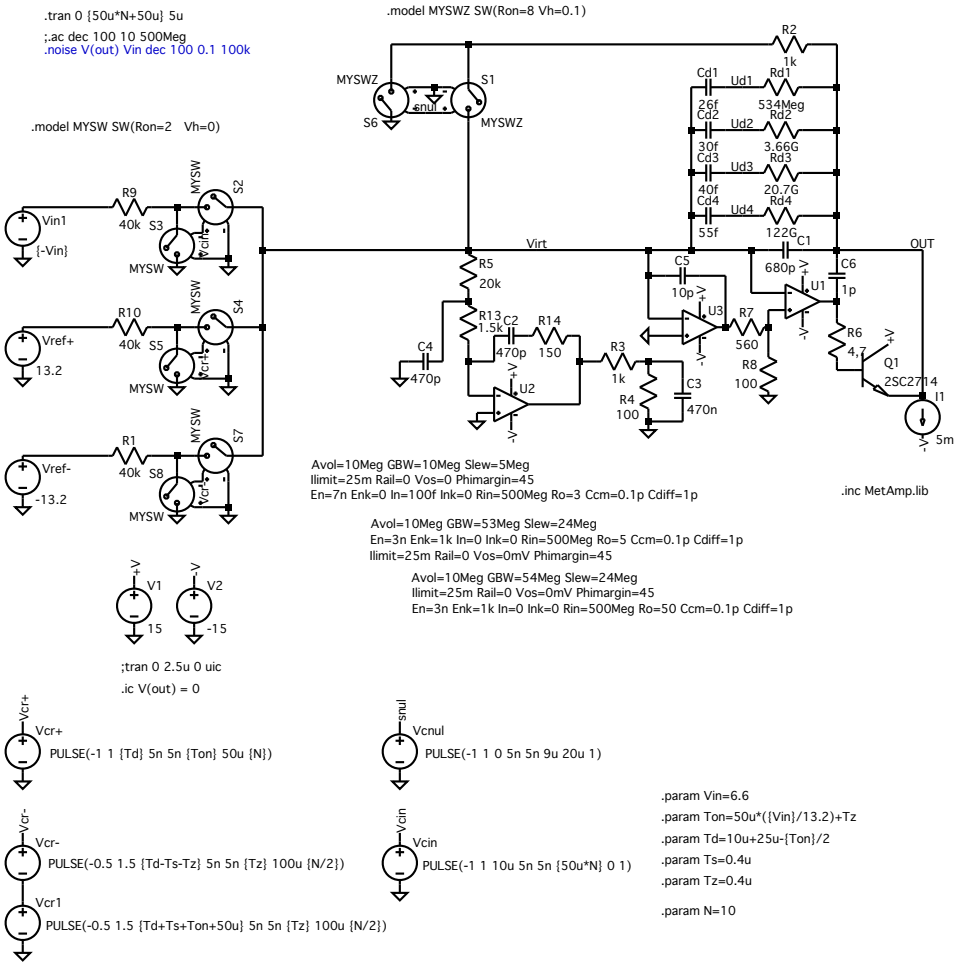

Figure 16: LTspice simulation of capacitor DA effect using PWM feedback modulation.

charge is not directly a deviation from linearity because its magnitude depends on the input voltage, but it is the worst possible estimate, the actual linearity error will be smaller. The voltages on the individual parasitic capacitors can be easily displayed, the maximum value is for the largest possible voltage swing at the integrator output, which occurs for $I_{\mathrm{in}} = I_{\mathrm{ref}}/2 = 165\,\mu\mathrm{A}$. As can be seen from the simulation result on Fig. 17, these voltages do not depend on the number of cycles and are 1.05 V on C_{d1} and 180 mV on C_{d2}, the voltages on the other capacitors are negligible, so the total charge loss due to DA is 33 fC. This charge corresponds to a time of 0.13 ns at a current of $I_{\mathrm{in}} = 250\,\mu\mathrm{A}$, i.e. 6.5 ppb of the integration time of 20 ms and is therefore negligible. Evaluating the effect of DA in "classical" integration used in HP3458 is even more complicated, depending on the details of the conversion and the amplitude of the input voltage, but it is still possible to make a rough estimate. We assumed a cycle time of 10 µs, the switching time of the reference switch is 8.8 µs. In the worst case, it may be that there is a sawtooth voltage of only one polarity at the integrator output with an amplitude of 7.1 V, the mean value of this voltage is 3.55 V and for simplicity we assume that the run-up phase ends when this voltage is at the

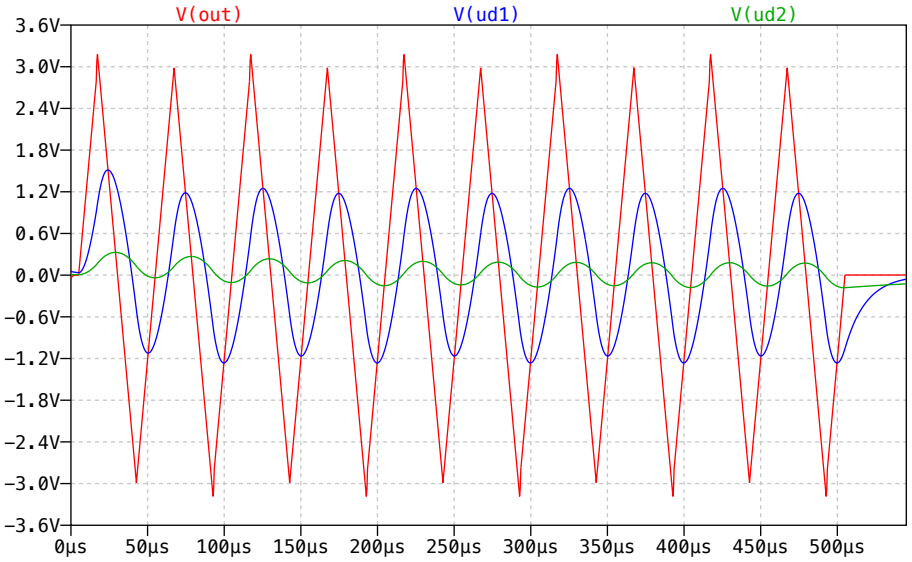

Figure 17: LTSpice simulation of capacitor DA effect using PWM feedback modulation.

output. Then, at longer integration times (tens of ms or more), all parasitic capacitors $C_{\mathrm{d1}} \cdots C_{\mathrm{d4}}$ will charge to this voltage and the total charge loss will be 530 fC. This corresponds to a time error of 2.1 ns at a current of 250 µA, i.e. 0.1 ppm at an integration time of 20 ms. This value is more than an order of magnitude larger than for PWM modulation and may no longer be completely negligible.

5.2 IADC system modelling using MATLAB

Fig. 18 shows a block diagram of a model programmed in MATLAB script, where based on the input signal, circuit parameters and the feedback switch modulation, all switch positions are determined and set during the IADC state calculation as described by the feedback modulation algorithm (block IADC Calculate). Using this information, IADC can also be simulated using a Simulink model (block IADC Simulate). All data can be stored and the post processing analysis can be performed (block IADC Post Processing).

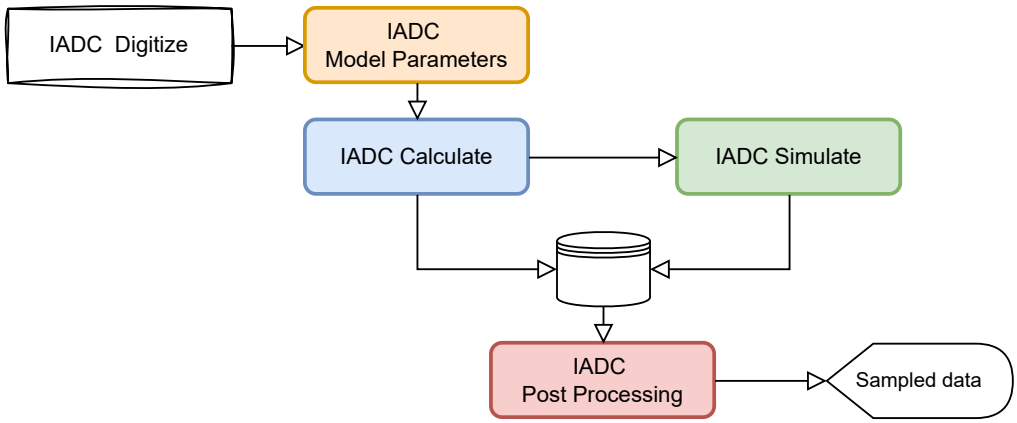

Figure 18: IADC model structure. Each coloured block represents a separate MATLAB script, that is called by IADC Digitize. The arrows denote call sequence and data flow.

For simulating a behaviour as a digitiser, the script **DigitizeIADC** defines digitising simulation parameters that can be adjusted and tested. It also calls all model scripts that comprise the whole IADC modelling/simulation package. The model/simulation script m-files are

- **DigitizeIADC**: The main script to be called with simulation parameters DEF parameter defined
- **IADCgenerateSimulationParameters**: This script takes DEF structure and creates initial IADC structure with model and simulation parameters, used by
- **IADCcalculate**: This procedure takes all IADC settings written in IADC structure and calculates the output in CalcIADC. It also returns IADC structure updated with actual Switches positions, Input voltages and Component values as time series, needed for Simulink simulation.
- **IADCsimulate**: This procedure takes updated IADC structure to provide all necessary information to perform Simulink simulation. Behind this, a Simulink file **IADCmodelS002.slx** is defined and is called from within the **IADCsimulate** script.

There are four data structures which carry modelling/simulation data and results:

```
1 DEF;              % parameters defining simulation run
2 IADC;             % all information defining simulation at each simulation step
```

```
CalcIADC;       % results from MATLAB calculation
SimIADC;        % results from Simulink simulation
```

DEF parameters define various aspects of the modelling/simulation and are defined as

```
%% define simulation parameters
% model settings
DEF.Simulate = false;               % enable Simulink simulation
DEF.alg  = "DualSlope";             % DualSlope | MultiSlope1
DEF.Nsim = 5;                       % number of samples
DEF.ta   = 20e-6;                   % aperture time
DEF.ts   = 20e-6;                   % sampling time
% input signal
DEF.Signal.Type = "sine";           % sine | DC
DEF.Signal.amplitude = 5.0;         % signal peak amplitude in V
DEF.Signal.frequency = 1000;        % signal frequency for 'sine' in Hz
DEF.Signal.phase     = -0.5073;     % signal phase in rad
%% IMPERFECTIONS
DEF.SimulateCintVoltageCoefficient  = false;
DEF.SimulateRinVoltageCoefficient   = false;
DEF.SimulateRinLoading              = false;
DEF.SimulateSwitchTimeJitter        = false;
DEF.SimulateSwitchChargeInjection   = false;
DEF.SimulateCOPAinputGroundBounce   = false;
% IADC integrating C voltage coefficient
DEF.Cvc = [ 1e-6 0];                % relative linear and quadratic term
% IADC Rin voltage coefficient
DEF.Rin.vc = [-1e-6 0];             % relative linear and quadratic term
% IADC Rin temperature drift due to load
DEF.Rin.alpha = 0.1e-6;             % temp. coef. in 1/ºC
DEF.Rin.heatCapacity = 4e-7;        % heat capacity in J/K, typ 4e-7 for 1206 SMD
DEF.Rin.thermalResistance = 77;     % thermal resistance in ºC/W, typ 77 for 1206 SMD
% IADC reference voltages and resistances
DEF.ref.Rp = 10000;                 % resistance in ohms
DEF.ref.Rm = 10000;                 % resistance in ohms
DEF.ref.Vp =  12;                   % voltage in V
DEF.ref.Vm = -12;                   % voltage in V
% time jitter
DEF.tj.SWin = 1e-12;                % time jitter in seconds
DEF.tj.SWrp = 1e-12;
DEF.tj.SWrm = 1e-12;
DEF.tj.distribution = 'normal';     % 'normal' | 'rectangular'
% Switch Charge Injection
DEF.Switches.CISWin.up   =  2e-12;  % charge injection in coulomb
DEF.Switches.CISWin.down = -2e-12;  % charge injection in coulomb
DEF.Switches.CISWrp.up   =  2e-12;  % charge injection in coulomb
DEF.Switches.CISWrp.down = -2e-12;  % charge injection in coulomb
DEF.Switches.CISWrm.up   =  2e-12;  % charge injection in coulomb
DEF.Switches.CISWrm.down = -2e-12;  % charge injection in coulomb
% Ground Bounce of integrator COPA
DEF.GroundBounce.Vup   = 2e-3;      % voltage level in V
DEF.GroundBounce.Vdown = -2.5e-3;   % voltage level in V
```

5.2.1 Capacitor Voltage Coefficient

Enabled by:

```
DEF.SimulateCintVoltageCoefficient = true;
```

Defined by:

```
DEF.Cvc = [ 1e-6 0];                % relative linear and quadratic term
```

The capacitance of a capacitor is typically considered independent of the input voltage under ideal conditions, meaning that the capacitance C, which is defined as the charge Q stored per unit voltage V (i.e., $C = \frac{Q}{V}$), remains constant. However, in real-world capacitors, there can be capacitance dependence on the input voltage, especially in non-ideal capacitors. Some capacitors, particularly those that use nonlinear dielectric materials (e.g., ceramic capacitors), can show significant capacitance variation with voltage. In these cases, the dielectric constant changes with the applied electric field, resulting in voltage-dependent capacitance. NP0 ceramic capacitors are linear and have stable capacitance independent of the applied voltage. Their dielectric material (typically titanium dioxide) is paraelectric, meaning it does not experience changes in its dielectric constant with voltage, leading to very stable capacitance. They are commonly used in high-precision and high-frequency applications due to their low losses and stable electrical properties. However, recent measurements using a special capacitance bridge have shown there are small nonlinearities that can be modelled with linear and quadratic term The voltage dependence of capacitance can be mathematically described by expanding the capacitance C as a function of the applied voltage V. If the voltage coefficient is small, the change in resistance can be approximated by a linear relationship:

$$C(V) = C_0 \cdot (1 + \alpha_C \cdot V + \beta_C \cdot V^2), \tag{3}$$

where $C(V)$ is the capacitance at voltage V and C_0 is the resistance at a reference voltage (for our purposes 0 V).

5.2.2 Resistance Voltage Coefficient

Enabled by:

```
DEF.SimulateRinVoltageCoefficient = true;
```

Defined by:

```
DEF.Rin.vc = [-1e-6 0];             % relative linear and quadratic term
```

The voltage coefficient of resistance (α_V) is defined as the relative change in resistance per unit change in voltage. Mathematically, it can be expressed as:

$$\alpha_V = \frac{1}{R} \cdot \frac{dR}{dV}, \tag{4}$$

where R is the resistance of the resistor and $\frac{dR}{dV}$ is the derivative of resistance with respect to the voltage.

This coefficient is typically expressed in parts per million per volt (ppm/V) or percent per volt ($\%\,\mathrm{V}^{-1}$). A positive voltage coefficient means that the resistance increases with increasing voltage, while a negative coefficient means the resistance decreases.

The voltage dependence of resistance can be mathematically described by expanding the resistance R as a function of the applied voltage V. If the voltage coefficient is small, the change in resistance can be approximated by a linear relationship:

$$R(V) = R_0 \cdot (1 + \alpha_R \cdot V + \beta_R \cdot V^2), \tag{5}$$

where $R(V)$ is the resistance at voltage V and R_0 is the resistance at a reference voltage (for our purposes 0V).

In IADC, the voltage coefficient of resistors is crucial, as even small changes in resistance can significantly affect the performance of the circuit.

- **High Precision Circuits**: Resistors with low voltage coefficients are preferred to minimize the impact of voltage changes on circuit performance. Tracking resistor configurations shall be used as much as practical.
- **Temperature Dependence**: It's also worth noting that temperature can affect the voltage coefficient. As temperature changes, so can the resistance, further complicating the relationship between voltage and resistance.

5.2.3 Changes in Resistor Due to Loading in the Time Domain

Enabled by:

```
DEF.SimulateRinLoading = true;
```

Defined by:

```
DEF.Rin.alpha = 0.1e-6;              % temp. coef. in 1/ºC
DEF.Rin.heatCapacity = 4e-7;         % heat capacity in J/K, typ 4e-7 for 1206 SMD
DEF.Rin.thermalResistance = 77;      % thermal resistance in ºC/W, typ 77 for 1206 SMD
```

The heating of a resistor is governed by the principle of conservation of energy, which can be expressed as a differential equation relating the power dissipated in the resistor to the heat lost to the environment. This can be reformulated using the concept of thermal resistance:

$$C\frac{dT}{dt} = I^2R - \frac{(T - T_{\text{ambient}})}{\theta_{\text{th}}} \tag{6}$$

Where:

- C is the heat capacity of the resistor.
- T is the temperature of the resistor.
- t is time.
- I is the current through the resistor.
- R is the resistance.
- θ_{th} is the thermal resistance of the resistor, expressed in °C/W.
- T_{ambient} is the ambient temperature.

This equation describes how the temperature of a resistor changes over time due to the balance of power dissipation and heat loss to the surroundings.

To model the change in resistance over discrete time intervals, we can convert the differential equation (6) into a difference equation. By considering small time steps Δt, the equation becomes:

$$C\frac{T_{n+1}-T_n}{\Delta t} = I_n^2 R_n - \frac{(T_n - T_{\text{ambient}})}{\theta_{\text{th}}} \tag{7}$$

Solving for T_{n+1}, we have:

$$T_{n+1} = T_n + \frac{\Delta t}{C}\left(I_n^2 R_n - \frac{(T_n - T_{\text{ambient}})}{\theta_{\text{th}}}\right) \tag{8}$$

Using the relationship between resistance and temperature, the resistance at the next time step is:

$$R_{n+1} = R_0 \times \left(1 + \alpha\left(T_n + \frac{\Delta t}{C}\left(I_n^2 R_n - \frac{(T_n - T_{\text{ambient}})}{\theta_{\text{th}}}\right) - T_0\right)\right) \tag{9}$$

This difference equation allows us to simulate how the resistance of a resistor changes over time in response to the current flowing through it.

Example, calculated in MATLAB using 9, is given for the following constants, where an initial temperature is 25 °C and no current is flowing through the resistor

- $C = 0.5$ J/K.
- $I = 10$ mA, dropping to 5 mA at 300 s.
- $R = 1300\,\Omega$.
- $\theta_{\text{th}} = 77$ K/W.
- $\alpha = 10$ ppm/°C.

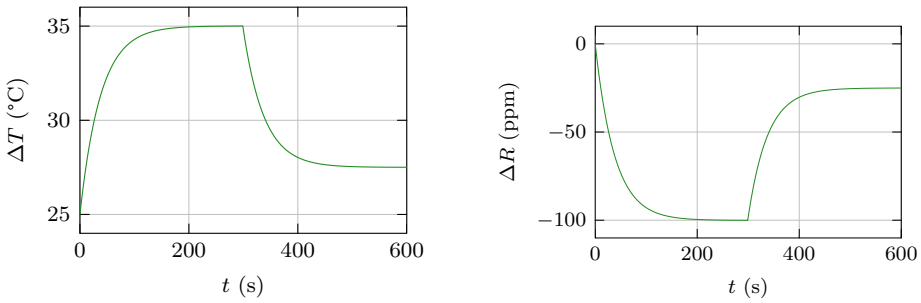

Figure 19: Example calculation of temperature (left) and resistance change (right) using eq. 9.

5.2.4 Switch time jitter

Enabled by:

```
DEF.SimulateSwitchTimeJitter = true;
```

Defined by:

```
DEF.tj.SWin = 1e-12;                % time jitter in seconds
DEF.tj.SWrp = 1e-12;
DEF.tj.SWrm = 1e-12;
DEF.tj.distribution = 'normal';     % 'normal' | 'rectangular'
```

The switch jitter is caused by non-ideal timing signals operating the switches and internal delays within the switches. They are modelled using a random number generator within the defined time jitter parameter for each switch separately. Either normal distribution or uniform distribution can be selected for the generation of random values. As this is a stochastic process, simulation will return different result with each run. Typically, one wants to observe the relation between switch jitter amplitude and output noise of the IADC converter.

5.2.5 Switch charge injection

Enabled by:

```
DEF.SimulateSwitchChargeInjection = true;
```

Defined by:

```
DEF.Switches.CISWin.up   =  2e-12;  % charge injection in coulomb
DEF.Switches.CISWin.down = -2e-12;  % charge injection in coulomb
DEF.Switches.CISWrp.up   =  2e-12;  % charge injection in coulomb
DEF.Switches.CISWrp.down = -2e-12;  % charge injection in coulomb
DEF.Switches.CISWrm.up   =  2e-12;  % charge injection in coulomb
DEF.Switches.CISWrm.down = -2e-12;  % charge injection in coulomb
```

Charge injection in CMOS switches refers to an undesired phenomenon that occurs when the MOSFET transistors used in CMOS circuits (specifically in switches) transfer or inject charge into the circuit during the switching process. This effect primarily arises due to the capacitance and switching behavior of the MOSFET's gate. A MOSFET switch is controlled by a gate voltage. When the gate voltage switches from "on" to "off" (or vice versa), the channel conducting current is either created or removed. During this transition, charges stored in the MOSFET's gate capacitance and other parasitic capacitances are released into the adjacent nodes (source or drain).

Charge Injection Process:

- Gate Capacitance: When the MOSFET is turned on, the gate is charged with a certain amount of charge due to its capacitance. When the gate voltage is switched, this charge is redistributed.
- Switching Transition: During the transition from on to off, the charge stored in the gate and channel (in the form of channel charge) gets injected into the source or drain. Since the switch is part of a circuit, this injected charge disturbs the signal or voltage level, causing inaccuracies.
- Effect on Signal: The injected charge can lead to voltage spikes or offsets at the output, degrading the performance of precision analog circuits like sample-and-hold circuits, analog-to-digital converters (ADCs), and digital-to-analog converters (DACs). This is particularly problematic when dealing with small voltage signals, where even a small injected charge can cause significant errors.

Charge injection is a critical issue in precision analog circuits, and minimizing its effects is a key design challenge for circuit engineers.

In Matlab script, charge injection is modelled by adding the following voltage to the integrator output:

$$U_{\mathrm{ci}} = \frac{Q}{C} = \frac{\Delta i \Delta t}{C} \tag{10}$$

In Simulink, it is simulated by injecting Δi into the integrator summing node for time Δt.

In both cases, the current is a constant value during the internal IADC clock period Δt, given by

$$\Delta i = \frac{Q}{\Delta t} \tag{11}$$

The values for charge injection simulation are taken from findings in Sec. 5.1.1.

5.2.6 COPA input ground bounce

Enabled by:

```
DEF.SimulateCOPAinputGroundBounce = true;
```

Defined by:

```
DEF.GroundBounce.Vup   = 2e-3;      % voltage level in V
DEF.GroundBounce.Vdown = -2.5e-3;   % voltage level in V
```

Composite amplifier experiences a ground voltage bounce when the current in capacitor changes rapidly during the switch operation. This can be nicely observed in LTSpice simulation circuit on virtual ground signal `Virt` (COPA circuit and LTSpice simulation scheme provided by John Pickering, see [12]).

In MATLAB script, ground bounce is modelled by adding the ground bounce voltage to the integrator output. In Simulink, it is simulated by applying the ground bounce voltage to the integrating operational amplifier positive node. As both simulations can apply this voltage over a predefined clock time period, set to 20 ns, the voltage level is chosen to represent the same voltage-time area as obtained from the LTSpice COPA simulation.

5.2.7 Simulink circuit

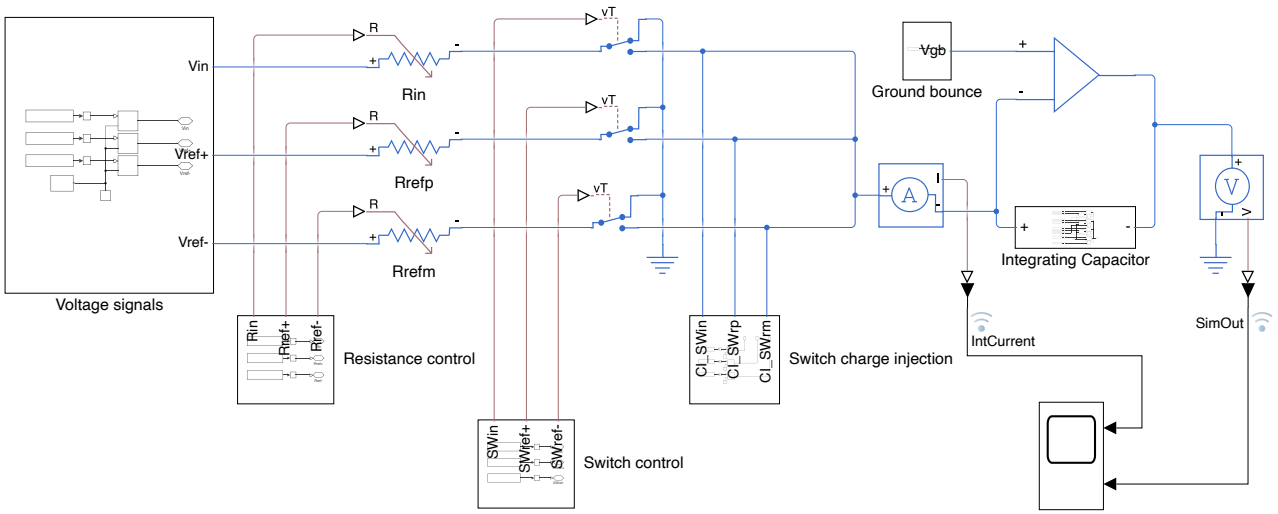

Figure 20: IADC top level Simulink model.

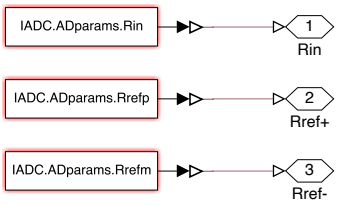

(a) Resistance control

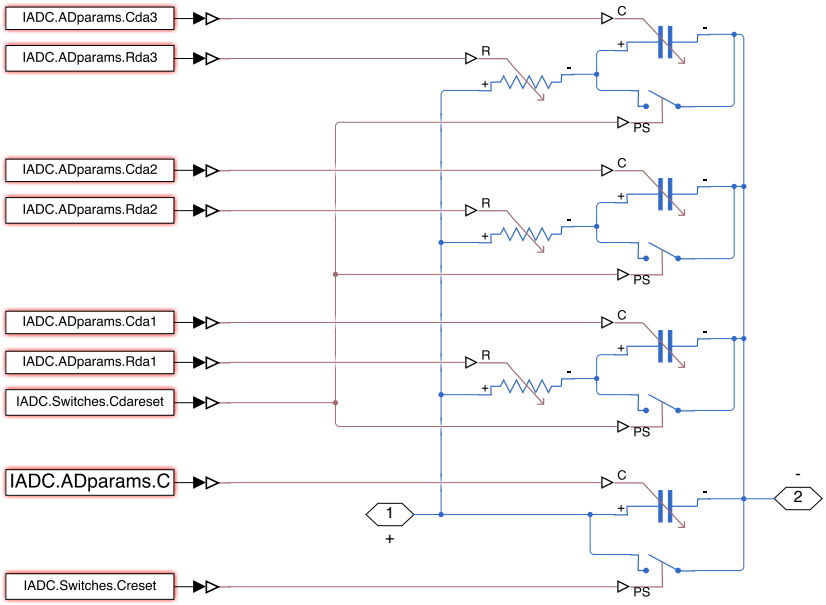

(b) Integrating Capacitor

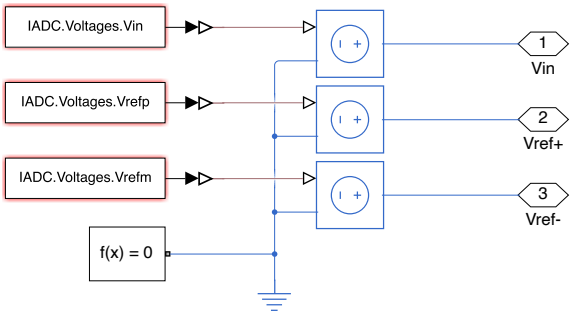

(a) Voltage signals

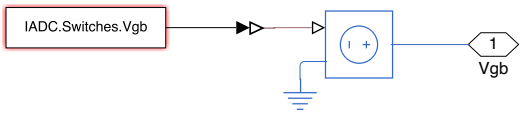

(b) Ground bounce

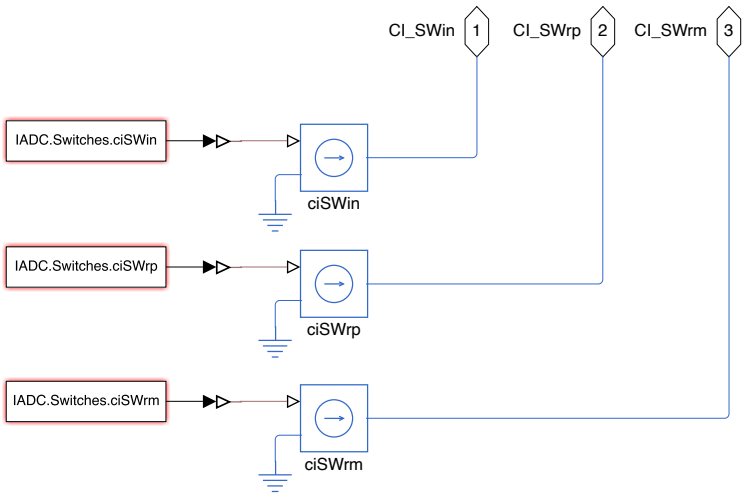

Figure 23: Switch charge injection

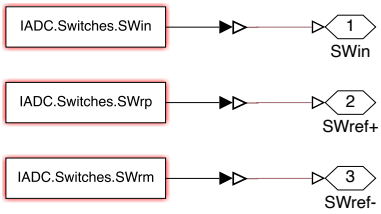

Figure 24: Switch control

5.3 IADC simulation results

The IADC simulation was performed using two feedback strategies, Dual Slope and Multi Slope. The integrating capacitor value was set to $C_{\text{int}} = 330\,\text{pF}$, while input and reference resistors were set to $10\,\text{k}\Omega$. Input signal was varied between $-10\,\text{V}$ and $10\,\text{V}$ and reference voltages were set at $\pm 12\,\text{V}$.

Using Dual Slope, the charge on integrating capacitor was first accumulated by input signal and then the number of clock periods were counted for its discharge using either $V_{\text{ref}+}$ or $V_{\text{ref}-}$ signal.

Using MultiSlope, while the charge was accumulating by input signal, it was simultaneously discharged by either $V_{\text{ref}+}$ or $V_{\text{ref}-}$ signal in a way that equal amount of switch transitions were used for both positive and negative reference in one conversion. At the end of aperture time, the remaining charge on the integrating capacitor was measured by SAR ADC. This way, the sampling time was set equal to aperture time.

The IADC simulations were performed for DC and AC input signals. Only the parameter of interest (non-ideal component behaviour) was varied, while the other parameters non-ideal behaviours were not simulated. This way, the various effects were isolated to be studied independently.

The DC simulation was performed by varying input signal from $-10\,\text{V}$ to $10\,\text{V}$ and observing IADC converted result. The aperture time was set to $50\,\mu\text{s}$.

The AC simulation was performed by applying a sine wave input signal with an amplitude of $10\,\text{V}$ and frequency of $2\,\text{kHz}$, except as noted along the result reported. A coherent sampling was simulated and the resulting converted record was evaluated using FFT with no windowing and observing harmonic content.

The simulator was checked for residual (numeric) errors with no component errors and consistently calculated conversions of input signal with less than $1 \times 10^{-14}\,\text{V/V}$ error. All DC linearity results were de-trended, i.e. the linear fit trend line was subtracted from the resulting deviation from the input DC voltage.

5.3.1 Resistor imperfections

Input resistor voltage coefficient DC simulation (Fig. 25) These results are independent of the feedback implementation. α_V corresponds to linear term and β_V corresponds to quadratic term (see Chapter 5.2.2). The aperture time was set to 50 µs.

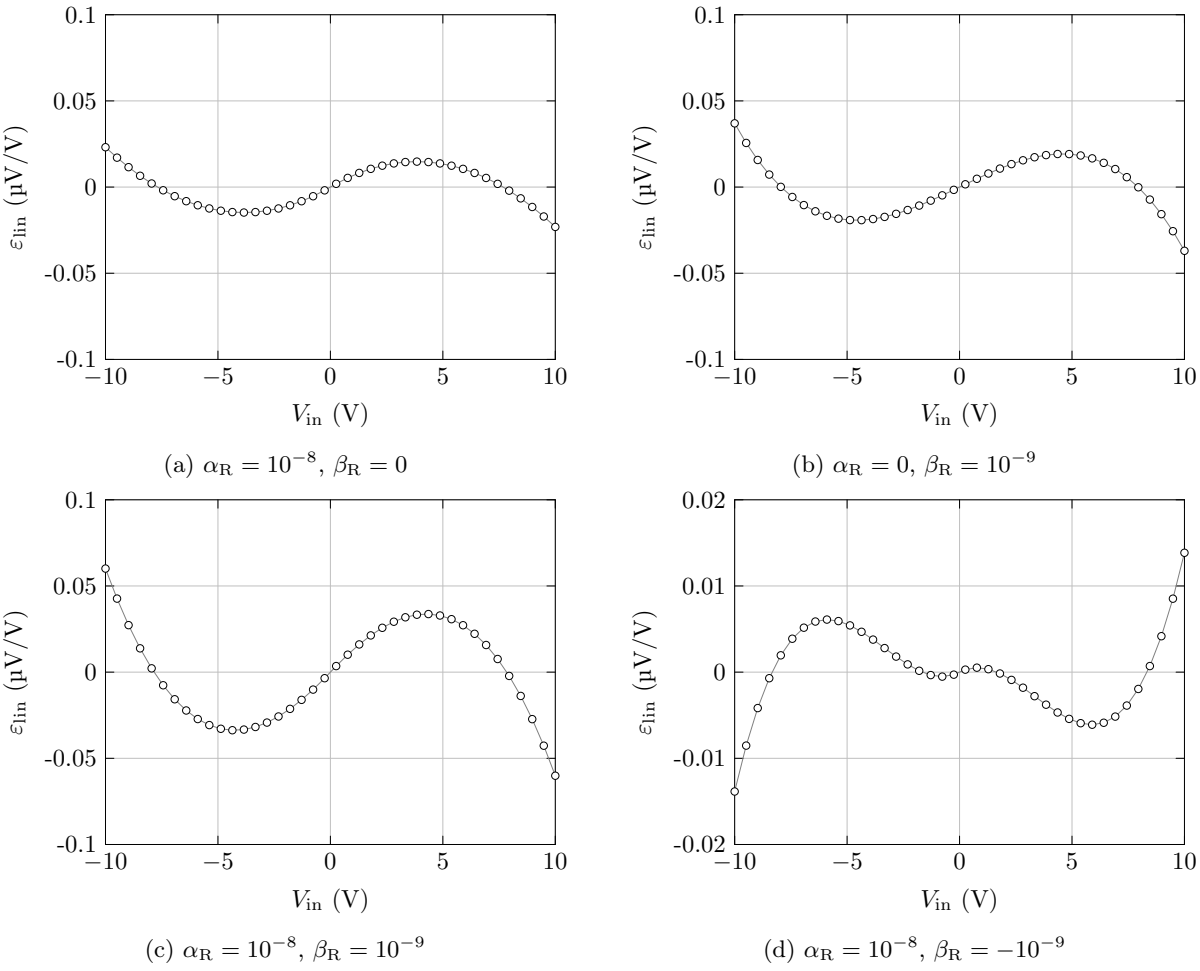

(a) $\alpha_R = 10^{-8}$, $\beta_R = 0$

(b) $\alpha_R = 0$, $\beta_R = 10^{-9}$

(c) $\alpha_R = 10^{-8}$, $\beta_R = 10^{-9}$

(d) $\alpha_R = 10^{-8}$, $\beta_R = -10^{-9}$

Figure 25: DC simulation linearity results for R_{in} voltage coefficient.

Input resistor voltage coefficient AC simulation (Fig. 26) These results are independent of the feedback implementation. The aperture time was set to 20 µs. 200 samples were calculated to produce each of spectrum plots.

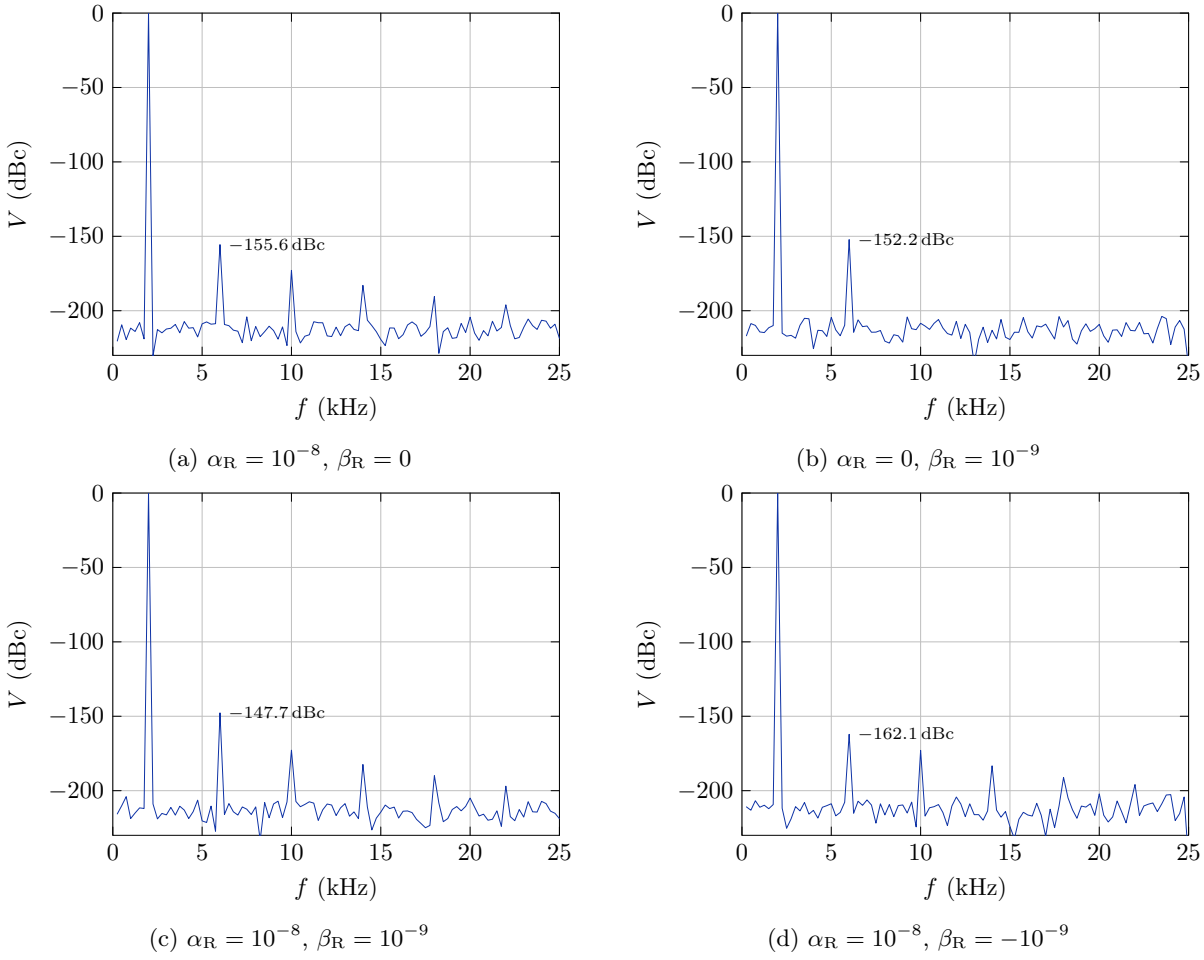

(a) $\alpha_R = 10^{-8}$, $\beta_R = 0$

(b) $\alpha_R = 0$, $\beta_R = 10^{-9}$

(c) $\alpha_R = 10^{-8}$, $\beta_R = 10^{-9}$

(d) $\alpha_R = 10^{-8}$, $\beta_R = -10^{-9}$

Figure 26: AC simulation linearity results for R_{in} voltage coefficient.

These results show that resistor voltage coefficients can impact IADC linearity. Referring to resistor non-linearity measurements performed [11], β_R is higher than the one used in this simulation for Metal Foil (Network) resistors ($\beta_R = 2.33(5) \times 10^{-8}$) and Carbon Film resistors ($\beta_R = 3.04(9) \times 10^{-9}$).

Input resistor loading DC simulation (Fig. 27) These results are independent of the feedback implementation. Input resistor parameters are: thermal resistance $\theta_{\text{th}} = 77\,°\text{C/W}$, resistor heat capacity $C = 4 \times 10^{-7}\,\text{J/K}$ and resistor temperature coefficient $\alpha = 0.1\,\mu\Omega/\Omega/\text{K}$. In DC mode, the heat capacity does not change results. Thermal resistance was varied to simulate its dependence on linearity. The aperture time was set to 50 µs.

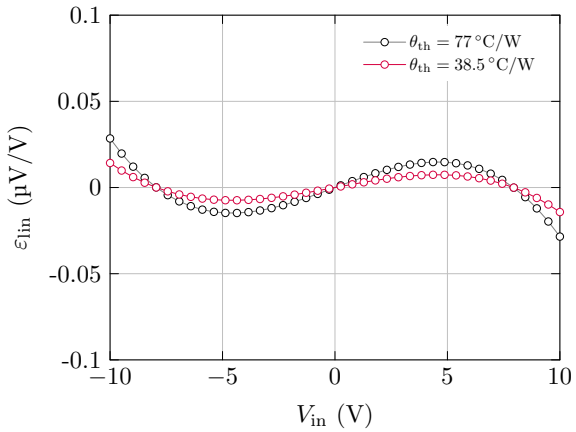

Figure 27: DC simulation linearity results for R_{in} loading effects.

Input resistor voltage coefficient AC simulation (Fig. 28) These results are independent of the feedback implementation. Heat capacity was varied to simulate its dependence on linearity, $t_{\text{a}} = 20\,\mu\text{s}$.

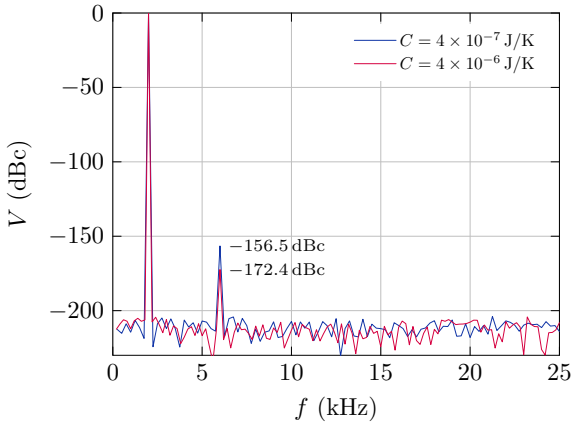

(a) $f = 2\,\text{kHz}$, $N = 200$.

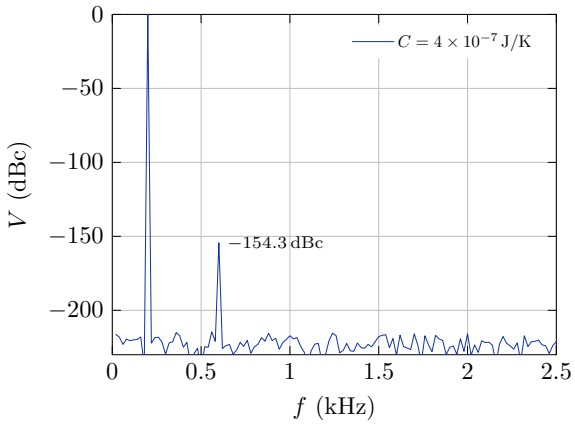

(b) $f = 200\,\text{Hz}$, $N = 2000$.

Figure 28: AC simulation linearity results for R_{in} loading effects. N stands for the number of samples.

These results show that resistor loading effect can be reduced by decreasing thermal resistance (DC and AC) and by increasing heat capacity. Plotted results were obtained for typical 1206 SMD chip resistor but with extremely low temperature coefficient of $\alpha = 0.1\,\mu\Omega/\Omega/\text{K}$. Therefore, this effect can reasonably contribute to IADC linearity.

5.3.2 Capacitor imperfections

Integrating capacitor voltage coefficient DC simulation (Fig. 29) These results clearly dependent of the feedback implementation. α_V corresponds to linear term and β_V corresponds to quadratic term (see Chapter 5.2.1). The aperture time was set to 50 µs.

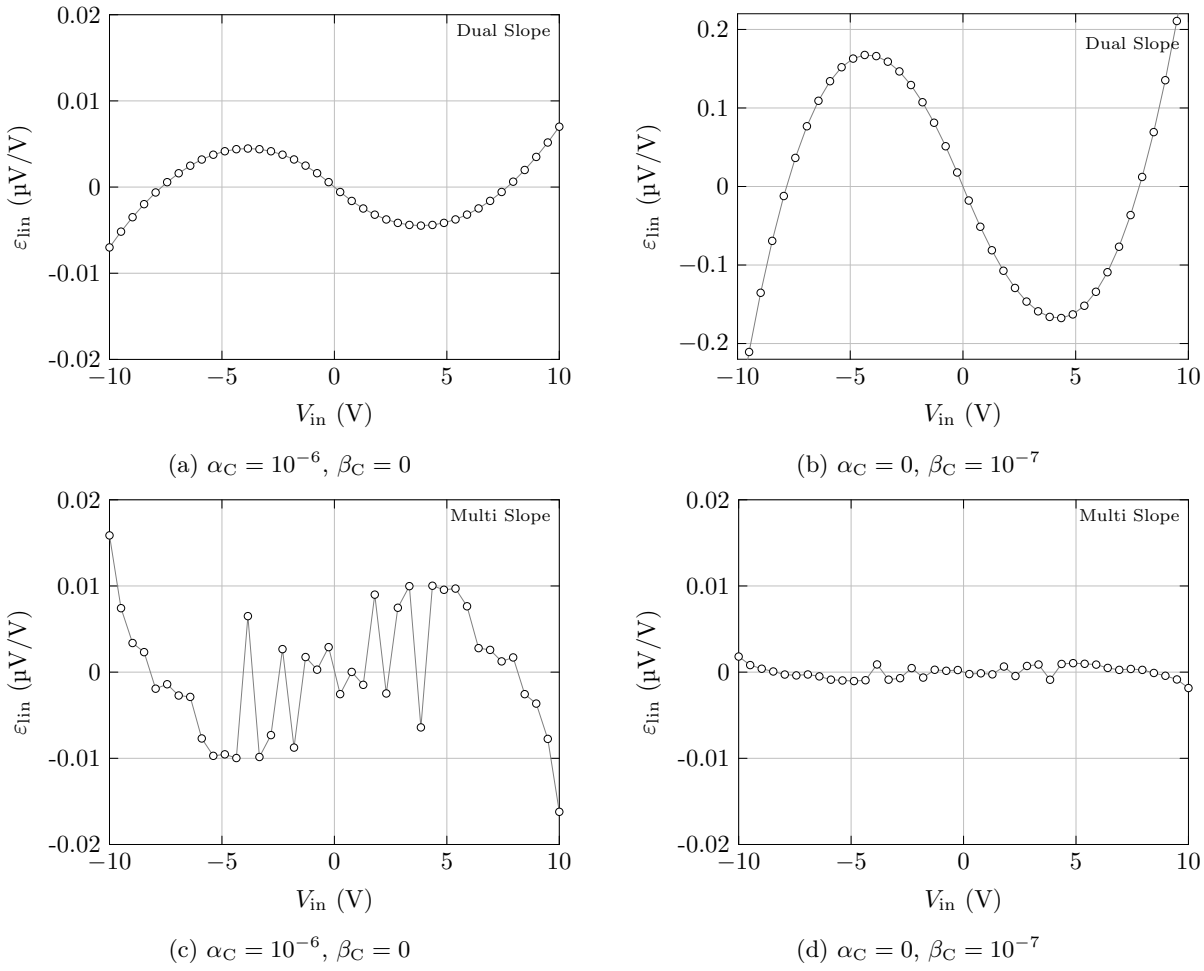

(a) $\alpha_C = 10^{-6}$, $\beta_C = 0$

(b) $\alpha_C = 0$, $\beta_C = 10^{-7}$

(c) $\alpha_C = 10^{-6}$, $\beta_C = 0$

(d) $\alpha_C = 0$, $\beta_C = 10^{-7}$

Figure 29: DC simulation linearity results for C_{int} voltage coefficient.

Integrating capacitor voltage coefficient AC simulation (Fig. 30) Here only results for Multi Slope feedback implementation are shown. The aperture time was set to 20 µs. 200 samples were calculated to produce each of spectrum plots.

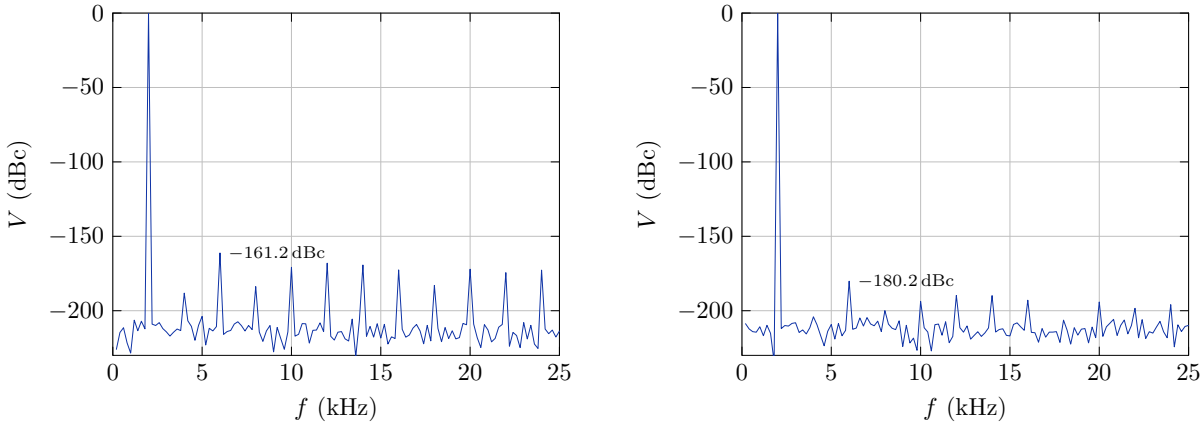

(a) $\alpha_C = 10^{-6}$, $\beta_C = 0$

(b) $\alpha_C = 0$, $\beta_C = 10^{-7}$

Figure 30: AC simulation linearity results for C_{int} voltage coefficient.

This results show strong dependence on the feedback implementation, where the Dual Slope returns smooth response, which is extremely strong for quadratic term β_C, while the Multi Slope strongly depends on the switch timings that govern capacitor charge for linear term α_C and becomes highly attenuated for quadratic term β_C. While the results show that the effect is small, it is not predictable and needs to be tested with the actual feedback implementation used. Further analysis of the actual mechanism could prove beneficial to better designing the feedback implementation to mitigate integrating capacitor voltage coefficient impact on linearity. As well, the capacitor α_C and β_C shall be measured across the frequency range of interest.

Integrating capacitor dielectric absorption DC simulation (Fig. 31) While simulation can cover multiple loops of dielectric absorption RC elements, only one loop was used to produce these results.

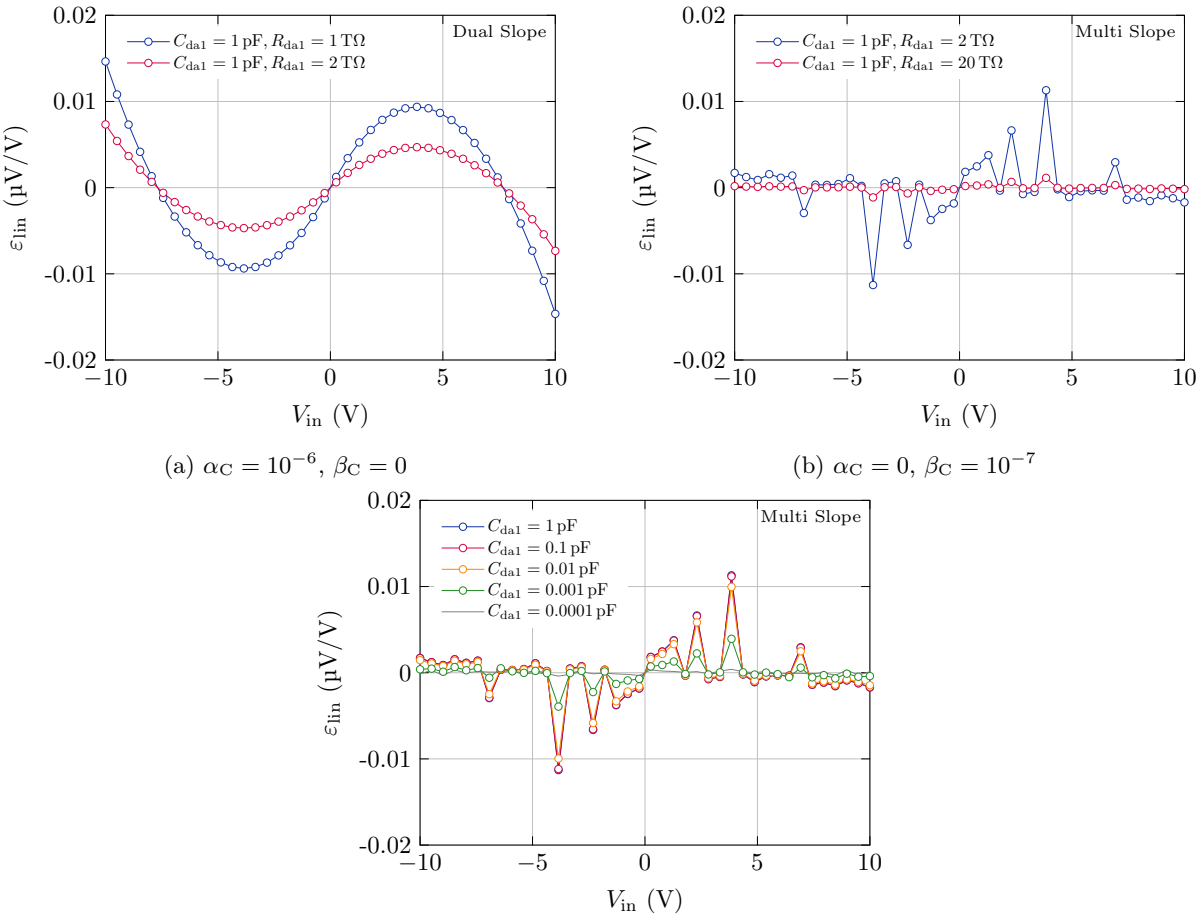

(a) $\alpha_C = 10^{-6}$, $\beta_C = 0$

(b) $\alpha_C = 0$, $\beta_C = 10^{-7}$

(c) $R_{da1} = 20\,\text{G}\Omega$

Figure 31: DC simulation linearity results for C_{int} dielectric absorption.

This results again show strong dependence on the feedback implementation, where the Dual Slope performing far worse than Multi Slope. Further, Multi Slope feedback implementation does not show smooth response as a function of input voltage, indicating that the effect depends on the actual switch timing ratios and charge accumulated on the integrating capacitor. Further analysis of the actual mechanism could prove beneficial to better design the feedback implementation to mitigate integrating capacitor dielectric absorption impact on linearity.

Figs. 31 (a) and (b) clearly show proportional impact of R_{da1} value on linearity. However, Fig. 31 (c) shows an unexpected observation, where it was found that the impact on linearity is almost negligible for $1\,\text{pF} \geq C_{da1} \geq 0.01\,\text{pF}$ whereas the effect is significant for lower values. The linearity deterioration is not proportional to the capacitance value, however.

These results indicate a need to measure short term integrating capacitor dielectric absorption, that can affect the IADC converter linearity, or to design a feedback that effectively cancels these effects.

5.3.3 Switches' charge injection

The accurate simulation of charge injection in switching elements within the Integrating ADC (IADC) model presents significant challenges. The phenomenon is highly dependent on the specific physical characteristics of the switches, the drive circuitry, and the dynamic operating conditions, which are difficult to model with sufficient reliability in the proposed high-level MATLAB simulation concept, which was designed to account only for a simple fixed charge injection at any switch toggle (separate for on and off), without further interactions and dependencies with the actual IADC circuit. A reliable simulation was provided through LTSpice modelling, as described and reported in Chapter 5.1.

Additionally, rather than relying on potentially misleading simulation data, this effect was investigated through direct measurement. Dedicated demo integrator circuits were developed in parallel to this modelling activity to physically characterize these imperfections. The results of these measurements, which provide a more reliable assessment of switch charge injection impact, are detailed in the Deliverable D4 report [12].

5.3.4 COPA imperfections

Similarly to switch charge injection, the non-ideal behaviors of the Composite Operational Amplifier (COPA), such as dynamic settling errors, transient thermal effects, and complex interactions with the feedback network, proved too complex for reliable representation within the scope of the current developed MATLAB simulation framework. "Again, these imperfections only dealt with the impact of a fixed ground bounce on the accumulated charge, without a comprehensive physical model. Therefore, this effects are way more suitable to be studied within dedicated electronic circuit simulators (like LTSpice) or be measured on actual circuits. Comprehensive results using LTSpice are provided in Chapter 5.1.

To ensure the validity of the findings, these COPA imperfections were also evaluated experimentally using the aforementioned demo integrator circuits. The empirical data and analysis regarding COPA performance and its influence on the IADC linearity are comprehensively reported in Deliverable D4 [12]. The effect of virtual ground point disturbance is analysed in Chapter 5.1.1.

6 Novel ADC Topology

Based on the investigation of barriers and mitigations, a hybrid digitiser topology is proposed as the final design choice. This architecture combines the strengths of Integrating ADC (DC accuracy, linearity) and wideband SAR ADC (AC performance, bandwidth).

6.1 Proposed Hybrid Design

The proposed mixed design consists of a low frequency Integrating ADC (IADC) path and a wideband SAR ADC path sharing a common auto zero and auto calibration functionality. As described in Section 4.1 above the SAR ADC parallel operation mainly demonstrated limited or negligible linearity improvement. However, since promising DC linearity improvement using IADC based digitisers was recently demonstrated (to be included in True8Digit Deliverable D2 report), a parallel structure with four IADC is proposed instead.

This topology offers several key benefits:

- **Parallel Operation:** Enables concurrent capture of accurate low frequency content (IADC) and high frequency wideband content (SAR). This allows for signal integrity checks and cross-validation.
- **Accuracy Transfer:** The state-of-the-art DC accuracy of the IADC can be transferred to the wideband SAR ADC.
- **Phase Correction:** The Wideband SAR ADC can measure the input signal with high temporal precision, providing data to correct phase errors in the IADC caused by trigger and conversion logic delays.

6.2 Auto-Calibration Options

To maximize the performance of this hybrid topology, extensive auto-calibration options are integrated into the design:

1. **Traceability via External Reference:** The design includes a Reference block capable of using an external 10 V reference. Through auto-calibration routines, this provides a direct traceability path for the digitiser.
2. **IADC to SAR Calibration:** During self-calibration (or DC/low frequency calibration), the IADC acts as the primary standard to linearise and correct the SAR ADC path, transferring its superior DC linearity to the wideband path.
3. **Input Path Self-Calibration:** Following the extended IADC configuration analysis, an input self-calibration switch is proposed. This allows U_{ref} to be coupled to the input signal path to measure and correct for the matching between the input signal path and the inverted reference path.
4. **Reference Current Balancing:** An auxiliary circuit can be employed to self-calibrate the positive and negative reference currents (I_{ref+} and I_{ref-}). The residual sum of these currents (which should ideally be zero) can be amplified and digitised to generate a correction factor accounting for buffer drift and switch unbalance.

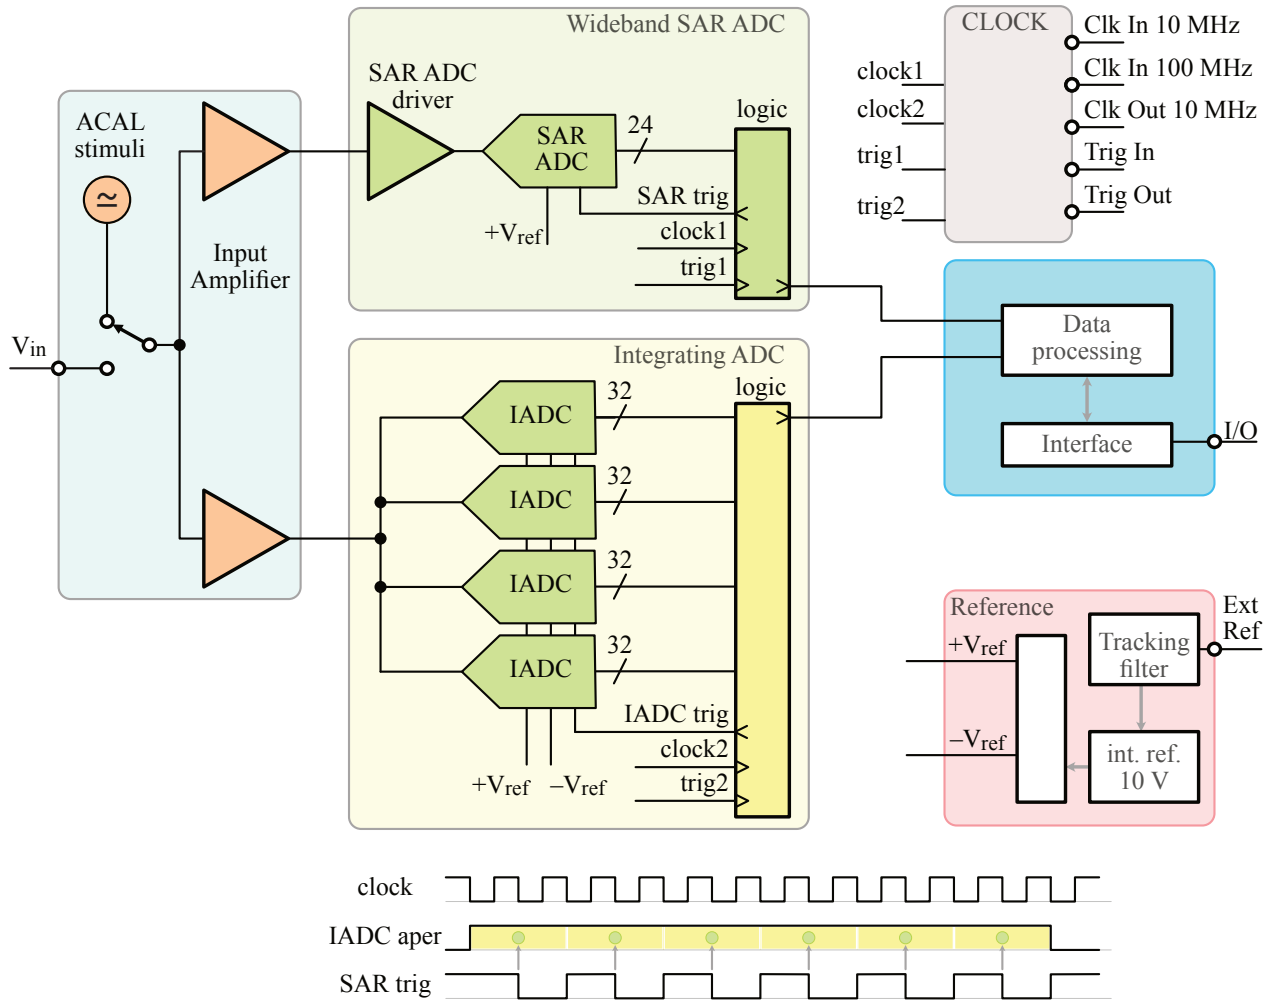

Figure 32: Digitiser using a hybrid topology, that assures DC accuracy and linearity (Integrating ADC path), as well as excellent AC performance, including flatness and low harmonic distortion (Wideband SAR ADC path). The input reversal for reducing pink noise and random walk is not shown for simplicity, but it shall be properly implemented.

References

[1] R. Lapuh, *Sampling with 3458A*. Left Right d.o.o., 2018.

[2] R. Lapuh, J. Kučera, J. Kováč, L. Palafox, and B. Voljč, "Digital Multimeters Voltage Sampling Performance Comparison," *Submitted for publication in IEEE TIM*, 2025.

[3] W. C. Goeke, "An 8 1/2-digit integrating analog-to-digital converter with 16-bit, 100,000-sample-per-second performance," *Hewlett-Packard Journal*, vol. 40, pp. 8–15, April 1989.

[4] Analog Devices, "Data Sheet AD4630-24/AD4632-24," tech. rep., Rev. B, 2025.

[5] R. Lapuh, B. Voljč, and M. Lindič, "Evaluation of Agilent 3458A Time Jitter Performance," *IEEE Trans. Instrum. Meas.*, pp. 1331–1335, 2015.

[6] N. Beev, "FGC 3.2 ANA200 Evaluation," tech. rep., CERN, 2020.

[7] W. Heinzel, A. Rüdiger, and R. Schilling, "Spectrum and spectral density estimation by the Discrete Fourier Transform (DFT), including a comprehensive list of window functions and some new flat-top windows.," tech. rep., Max-Planck-Institut für Gravitationsphysik, 2002. Also available as `http://pubman.mpdl.mpg.de/pubman/item/escidoc:152164:1/component/escidoc:152163/395068.pdf`.

[8] J. Schoukens, R. Pintelon, and G. Vandersteen, "A sinewave fitting procedure for characterizing data acquisition channels in the presence of time base distortion and time jitter," *IEEE Trans. Instrum. Meas*, vol. 46, no. 4, pp. 1005–1010, 1997.

[9] R. Lapuh, N. Beev, A. Belcher, and L. Palafox, "Digitizer barriers and mitigations," tech. rep., True8Digit, 2024.

[10] True8Digit, "Deliverable D3: Report on Metrological methods for characterisation of resistors, capacitors and switches," tech. rep., True8Digit, 2024. Available on https:https://zenodo.org/records/16410882.

[11] V. Janásek, R. Lapuh, L. Palafox, N. Beev, and B. Voljč, "A modified wheatstone bridge for high-precision characterisation of resistor nonlinearity," *To be submitted to IEEE Trans. Instrum. Meas*, 2026.

[12] True8Digit, "Deliverable D4: Report on two new composite operational amplifiers," tech. rep., True8Digit, 2026. Available on https:https://zenodo.org/records/18300606.

[13] J. Pickering, D. Georgakopoulos, J. Williams, and P. Wright, "Effect of a pwm feedback dac on the noise and linearity of a delta-sigma adc," in *5th IEE International Conference on ADDA 2005. Advanced A/D and D/A Conversion Techniques and their Applications*, January 2005.

[14] J. Pickering, D. Georgakopoulos, and J. Williams, "Theoretical limitations and physical realisation of a high resolution delta-sigma adc based on a pwm feedback dac," *The Mediterranean Journal of Electronics and Communications*, vol. 3, no. 2, pp. 7–11, 2007.

[15] J. Kuenen and G.C.M.Meijer, "Measurement of dielectric absorption of capacitors and analysis of its effects on vcos," *IEEE Trans. Instrum. Meas*, vol. 45, no. 1, pp. 89–97, 1996.

Appendix A Even-Order Distortion Cancellation Using Inverted Digitizer Channels

Digitizers and analog-to-digital converters (ADCs) exhibit nonlinearities that manifest as harmonic distortion when processing sinusoidal signals. In precision measurement systems, even-order harmonics are often particularly undesirable, as they indicate asymmetry in the transfer characteristic. This chapter analyzes a technique using two identical digitizers driven by signals of opposite polarity. By subtracting the digitized outputs, even-order harmonic distortion can be cancelled. The derivation is carried out using a sinusoidal representation, and the effect on noise is quantified.

Let the input signal be a pure sinusoid:

$$x(t) = A\sin(\omega t). \tag{12}$$

Each digitizer is assumed to be identical and memoryless, producing harmonic distortion that can be expressed as a sum of sinusoidal components:

$$y_1(t) = A\sin(\omega t) + \sum_{k=2}^{\infty} H_k \sin(k\omega t + \phi_k) + n_1(t), \tag{13}$$

where H_k and ϕ_k are the amplitude and phase of the k-th harmonic, and $n_1(t)$ is additive noise. The second digitizer is driven by the inverted input signal:

$$-x(t) = -A\sin(\omega t). \tag{14}$$

Inverting a sine wave is equivalent to a phase shift of π:

$$-A\sin(\omega t) = A\sin(\omega t + \pi). \tag{15}$$

Using the identity:

$$\sin(k(\theta + \pi)) = (-1)^k \sin(k\theta), \tag{16}$$

the output of the second digitizer becomes:

$$y_2(t) = -A\sin(\omega t) + \sum_{k=2}^{\infty} (-1)^k H_k \sin(k\omega t + \phi_k) + n_2(t). \tag{17}$$

The reconstructed signal is obtained by subtracting the two digitizer outputs and scaling by one half:

$$y(t) = \frac{y_1(t) - y_2(t)}{2}. \tag{18}$$

Substituting the expressions for $y_1(t)$ and $y_2(t)$ yields:

$$\begin{aligned} y(t) = {} & \frac{2A\sin(\omega t)}{2} \\ & + \frac{1}{2}\sum_{k=2}^{\infty} \left(1 - (-1)^k\right) H_k \sin(k\omega t + \phi_k) \\ & + \frac{n_1(t) - n_2(t)}{2}. \end{aligned} \tag{19}$$

The term $\left(1-(-1)^k\right)$ determines which harmonics remain. For even k: $(-1)^k = +1 \Rightarrow 1-(-1)^k = 0$. For odd k: $(-1)^k = -1 \Rightarrow 1-(-1)^k = 2$. Thus, all even-order harmonics cancel exactly, while odd-order harmonics remain:

$$y(t) = A\sin(\omega t) + \sum_{\substack{k\geq 3 \\ k \text{ odd}}} H_k \sin(k\omega t + \phi_k) + \frac{n_1(t) - n_2(t)}{2}. \tag{20}$$

Assuming the noise terms $n_1(t)$ and $n_2(t)$ are uncorrelated, zero-mean, and have equal variance σ^2, the variance of the noise component in the reconstructed signal is:

$$\mathrm{Var}\left(\frac{n_1 - n_2}{2}\right) = \frac{1}{4}\left(\sigma^2 + \sigma^2\right) = \frac{\sigma^2}{2}. \tag{21}$$

The resulting noise RMS is therefore:

$$\sigma_{\text{out}} = \frac{\sigma}{\sqrt{2}}. \tag{22}$$

This corresponds to a 3 dB improvement in signal-to-noise ratio relative to a single digitizer.

The differential technique provides two key benefits:

- Exact cancellation of even-order harmonic distortion for matched digitizers,
- Reduction of uncorrelated noise by a factor of $\sqrt{2}$.

In practice, the achievable cancellation is limited by gain mismatch, timing skew, reference noise correlation, and temperature-dependent drift. Nevertheless, this approach is highly effective in precision digitization and metrological applications.

By exploiting symmetry and signal inversion, even-order nonlinear distortion in digitizers can be eliminated using simple post-processing. The analysis presented here demonstrates that this result follows directly from sinusoidal identities and basic noise statistics, without requiring a polynomial nonlinearity model.