# Temperature-Robust Impedance Boosting and Chopper Stabilization for a Transformer-First Analog Front End: A Model-Based Study

Nithin Purushothama

nithinpurushothama@gmail.com • interactive model and full report: `https://nithinpuru.com/spa_afe/`

November 2026

## Abstract

The transformer-first analog front end of Atzeni *et al.* [1] achieves a noise efficiency factor (NEF) of 0.34 by using a discrete-time series-parallel capacitor transformer as its first gain stage, and recovers the input resistance lost to the bottom-plate parasitic capacitances with a calibrated precharge loop that raises it from 160 kΩ to 400 MΩ. This work presents a block-level behavioural model of that front end, with the transformer solved at charge level, validated against the published gain, noise floor, input resistance and input impedance. The model identifies two limitations. First, the 400 MΩ input resistance requires the parasitic charge to be restored to within approximately 0.01 %; after a single calibration it is therefore maintained only between −1.2 and 3.8 K around the calibration temperature, and at lower temperatures it becomes negative, which drives a capacitively coupled input to the supply rail. Second, below a few hundred hertz the input-referred noise is dominated by the 1/f noise of the amplifiers that follow the transformer. Six circuit techniques are proposed and evaluated in the same model. The principal one, a drift-tracking precharge (**P4**), implements the precharge capacitors as replicas of the parasitic capacitance and drives them from a unity-gain buffer of the transformer output. The buffer, which the recent SPA neural interface of the same group also uses [2], reduces the drift from 90 to 22 ppm/K; the replica capacitors reduce it further to 4 ppm/K, so that the input resistance remains above 100 MΩ over at least 92 K instead of 5 K (21 K with the buffer alone). Chopper stabilization within the transformer (**P2**) reduces the noise between 0.5 and 100 Hz from 0.37 to 0.093 $\mu V_{rms}$ and, combined with **P4**, requires no precharge hold. A delta transformer (**P6**), which samples relative to a predicted input, extends the input range from ±13 to ±250 mV at the cost of 3 dB of additional transformer noise. The remaining techniques add a margin and a temperature trigger to a binary-search calibration (**P1**), halve the precharge charge (**P3**) and lower the NEF to 0.25 through a larger transformer ratio enabled by **P4** (**P5**). All results are obtained from behavioural simulation under stated assumptions.

*Index Terms*: analog front end, biopotential acquisition, chopper stabilization, discrete-time amplifier, input-impedance boosting, series-parallel switched-capacitor transformer.

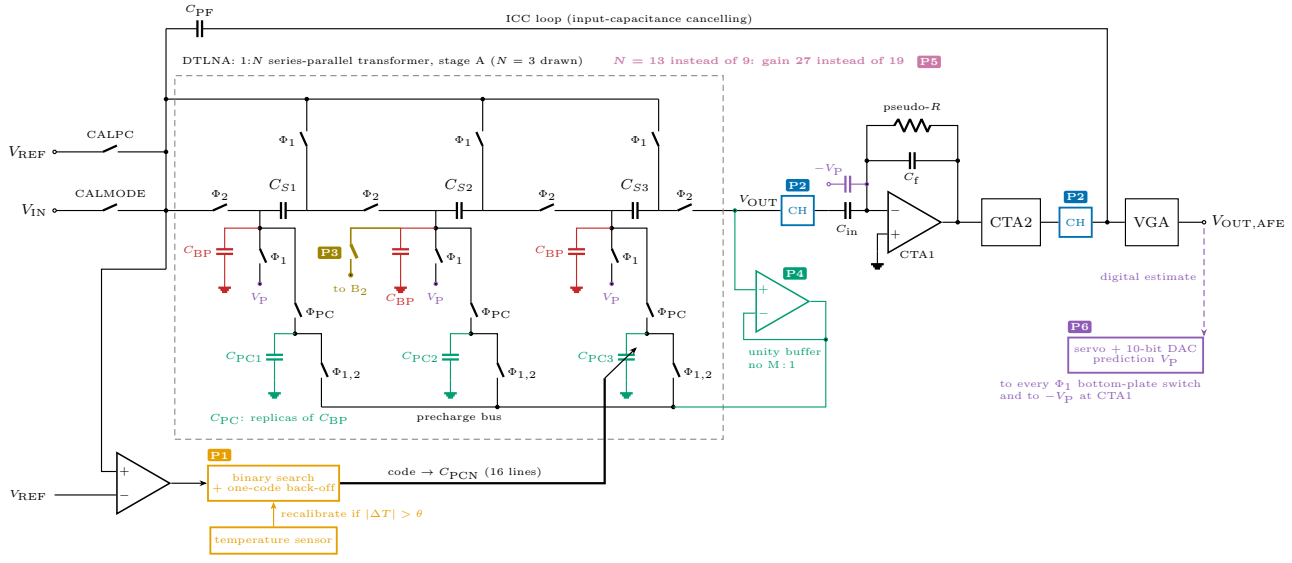


**Figure 1:** Proposed front end: the published circuit [1] (black; bottom-plate parasitic capacitances in red; one stage with $N = 3$ drawn) with the six techniques in colour. **P4** drives the precharge bus from a unity-gain buffer of the transformer output, taken ahead of the **P2** chopper. **P6** connects the sampling reference to a prediction $V_P$. **P1** modifies the calibration logic, **P3** adds one sharing switch per bottom plate (one drawn) and **P5** stacks 13 capacitors instead of 9.

**Table 1:** Proposed techniques, in order of significance. Improvements are model results relative to the published design.

| | Technique | Principle | Modelled improvement | Overhead |
|---|---|---|---|---|
| **P4** | Drift-tracking precharge | Precharge capacitors implemented as replicas of the bottom-plate parasitic capacitance, driven by a unity-gain buffer of the transformer output (buffer as in [2]) instead of CTA1 and the M:1 down-converter | Drift: 90 to 22 ppm/K with the buffer, 4 ppm/K with replicas; $R_{IN} > 100\,M\Omega$ over $\geq 92\,K$ instead of 5 K | Buffer bias current; replica matching |
| **P2** | Chopper stabilization within the transformer | Transformer output chopped at 2 kHz and demodulated after CTA2; with **P4** the precharge is derived ahead of the chopper, so no hold is required | Noise between 0.5 and 100 Hz reduced from 0.37 to $0.093\,\mu V_{rms}$ | Chopper switches; ripple filtering |
| **P6** | Delta transformer | Bottom plates sampled relative to a predicted input $V_P$ (10-bit DAC and servo); $V_P$ subtracted at the CTA1 input | Input range extended from $\pm 13$ to $\pm 250$ mV; no saturation with a $100\,mV_{pp}$ artifact | 3 dB additional noise; DAC |
| **P1** | Margin-preserving recalibration | One-code back-off and temperature-triggered recalibration, added to a binary-search calibration as in [2] | $R_{IN}$ never negative over 8 h; 20 instead of 489 comparisons with the binary search | Temperature sensor; 59 recalibrations in 8 h |
| **P3** | Charge recycling | Bottom plates shared with the corresponding plates of the other interleaved stage before the precharge | Precharge charge halved; drift reduced by a factor of 1.5 | $2N$ switches; one short clock phase |
| **P5** | Larger transformer ratio | $N = 13$ instead of 9 at the same total capacitance, enabled by **P4** | NEF reduced from 0.31 to 0.25; PEF from 0.084 to 0.053 | Input range reduced to 0.7 times; 9 % more power |

# 1 Introduction

Discrete-time front ends based on a series-parallel switched-capacitor transformer obtain passive voltage gain whose noise is set by $kT/C$ rather than by the transconductance efficiency of an input transistor [1, 3]. In [1], a 1:9 transformer provides a differential gain of 19 ahead of the first amplifier (CTA1) and enables an NEF of 0.34. The bottom-plate parasitic capacitances of the stacked capacitors, however, draw charge from the input in every cycle and would limit the input resistance to approximately 160 kΩ. An input-resistance boosting (IRB) loop restores this charge from precharge capacitors driven by CTA1 through an M:1 down-converter, and a foreground calibration trims one precharge capacitor to raise the input resistance to 400 MΩ.

This work examines the margin of that calibration and the origin of the remaining low-frequency noise by means of a behavioural model, and evaluates six circuit techniques in the same model. Sec. 2 describes the model and its validation, Sec. 3 the limitations it reveals, and Sec. 4 the proposed techniques. All results are obtained from simulation; no circuit was fabricated or measured.

# 2 Behavioural Model and Validation

The transformer is solved at charge level on every floating node of the capacitor stack, including the bottom-plate parasitic capacitance $C_{BP}$ of each capacitor and the interleaved second stage. The IRB loop is modelled as precharge capacitors charged from CTA1 through the M:1 down-converter, with finite loop gain and temperature coefficients. The calibration reproduces the comparator test on the floating input described in [1], and CTA1, CTA2 and the variable-gain amplifier (VGA) include white and 1/f noise fitted to Fig. 10(c) of [1]. The blocks are simulated together in a time-domain loop at the 1 MHz sampling clock. Prior to any modification, the model reproduces the transformer gain (25.58 dB versus 25.6 dB), the input resistance before and after calibration (158 kΩ and 390 MΩ versus 160 kΩ and 400 MΩ), the $kT/C$ noise floor (3.8 versus $4\,nV/\sqrt{Hz}$) and the measured $|Z_{IN}|$ of Fig. 10(b) of [1] within 0.06 decade rms (Fig. 2). Quantities not reported in [1], chiefly the open-loop gain of CTA1 (assumed 60 dB) and the temperature

**Table 2:** Published design and proposed techniques. Proposed values are at $N = 9$ except in the P5 rows. Drift results assume a 60 dB CTA1 open-loop gain and the temperature coefficients of Sec. 3.

| | **Published, measured** [1] | **Published, our model** | **Proposed, our model** | **Change** | **By** |
|---|---|---|---|---|---|
| *Input resistance* | | | | | |
| $R_{IN}$ after calibration | 400 MΩ (Sec. IV) | 390 MΩ | 389 MΩ | **unchanged** | |
| $R_{IN}$ 3 K below the calibration temperature | not reported | −290 MΩ (negative) | 434 MΩ | **remains positive** | P4 |
| Temperature span with $R_{IN} > 100$ MΩ after one calibration | not reported | 5.0 K (−1.2 to +3.8 K) | ≥ 92 K (−32 to > +60 K) | **≥ 18× wider** | P4 |
| Drift of the unrestored charge fraction $\delta$ | not reported | 90 ppm/K | 4.0 ppm/K | **22× lower** | P4 |
| Fraction of an 8-hour profile with negative $R_{IN}$ (single calibration) | not reported | 12.6 % | 0 % (never below 377 MΩ) | **never negative** | P4 |
| Comparisons per calibration (binary search, as in Chen et al., VLSI 2026) | linear search (Sec. III-E) | 489 (35 ms) | 20 (1.4 ms) | **24× fewer** | P1 |
| *Low-frequency noise* | | | | | |
| Input-referred noise, 0.5-100 Hz | not reported | 0.37 μVrms | 0.093 μVrms | **4.0× lower** | P2 |
| Input-referred noise, 1-40 Hz | not reported | 0.30 μVrms | 0.060 μVrms | **5.1× lower** | P2 |
| *Input range* | | | | | |
| DC-coupled input range (±0.25 V transformer output, assumed) | THD −54 dB at 10 mVpp | ±13 mV | ±250 mV (prediction DAC range) | **19× wider** | P6 |
| 100 mVpp, 2 Hz artifact with a 200 μV, 40 Hz signal | not reported | saturated 83 % of the time; 19 % of the signal recovered | no saturation; 100 % of the signal recovered | **signal recovered** | P6 |
| *Noise efficiency (optional P5)* | | | | | |
| NEF / PEF (10 kHz band) | 0.34 / 0.10 (Table I) | 0.31 / 0.084 | 0.25 / 0.053 with N = 13 | **1.2× / 1.6× lower** | P5 |
| Temperature span with $R_{IN} > 100$ MΩ at N = 13 | | 2.6 K (N = 13 without P4) | 53 K (−16 to +38 K) | **11× wider than published** | P5 P4 |
| *Costs* | | | | | |
| Transformer white-noise floor | 4 nV/$\sqrt{\text{Hz}}$ (Sec. IV) | 4.1 nV/$\sqrt{\text{Hz}}$ | 5.8 nV/$\sqrt{\text{Hz}}$ | +3 dB | P6 |
| Power | 370 nW (Sec. IV) | 370 nW | +33 nW for N = 13; P4 buffer and P6 DAC not modelled | increased | P4 P5 P6 |

coefficients, are assumptions, and the drift results scale with them.

## 3 Limitations of the Published Design

**Input resistance over temperature.** Let $\delta$ denote the fraction of the parasitic charge that is not restored by the precharge. The input resistance is then $R_{IN} = R_0/\delta$ with $R_0 \approx 46\,\text{k}\Omega$, so that 400 MΩ corresponds to $\delta \approx 0.01\,\%$. One calibration step changes $\delta$ by 0.023 %, which exceeds $\delta$ itself, and $\delta$ drifts by 90 ppm/K. The drift has two mechanisms: the closed-loop gain error of CTA1, which follows its open-loop gain (assumed −0.4 %/K), and the different temperature coefficients of the metal-oxide-metal (MOM) precharge capacitors and the junction-like parasitic capacitance (assumed 30 and 100 ppm/K, respectively).

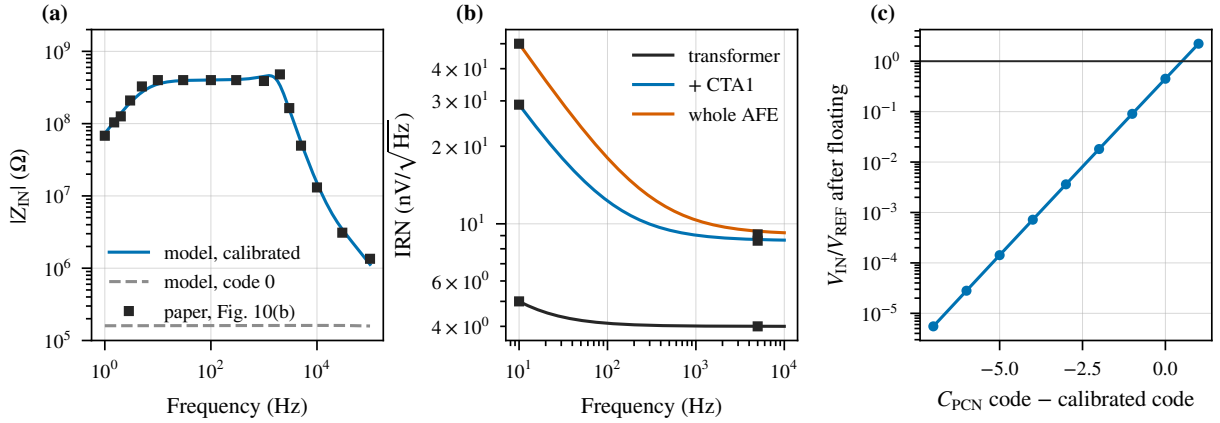


**Figure 2:** Model compared with the measurements of [1]. (a) $|Z_{IN}|$ after calibration and at code 0 (lines), with the measured points of Fig. 10(b) (squares). (b) Input-referred noise of the transformer, of the transformer with CTA1 and of the complete front end (lines), with the values read from Fig. 10(c) at 10 Hz and the reported noise floors (squares). (c) Calibration: input voltage relative to $V_{REF}$ after the floating interval, for codes around the calibrated one; the sign of the droop changes between two adjacent codes.

In the model, removing only the first leaves 22 ppm/K, removing only the second leaves 54 ppm/K, and removing both leaves 4 ppm/K. A temperature decrease over-restores the charge. After a single calibration, $R_{IN}$ remains positive and above 100 MΩ only between −1.2 and 3.8 K (Fig. 4); with a capacitively coupled electrode, a negative $R_{IN}$ causes the input to diverge toward the supply rail (Fig. 3(c)). Over a synthetic 8-hour skin-temperature profile, $R_{IN}$ is negative for 12.6 % of the time. The SPA-based neural interface of [2] reports a calibrated input impedance between 7.4 and 46.5 MΩ across 16 channels at 1 Hz, a spread consistent with a residual error that lies anywhere within one calibration step; no temperature dependence is reported.

**Low-frequency noise.** The transformer contributes little 1/f noise, whereas CTA1 contributes 64 % of the noise power in the 10 kHz bandwidth. Between 0.5 and 100 Hz, the 1/f noise of the amplifiers raises the input-referred noise to 0.37 $\mu V_{rms}$, four times the white-noise floor (Fig. 3). In electroencephalography (EEG) and electrocardiography (ECG) bands, the noise is therefore set by the amplifiers rather than by the transformer.

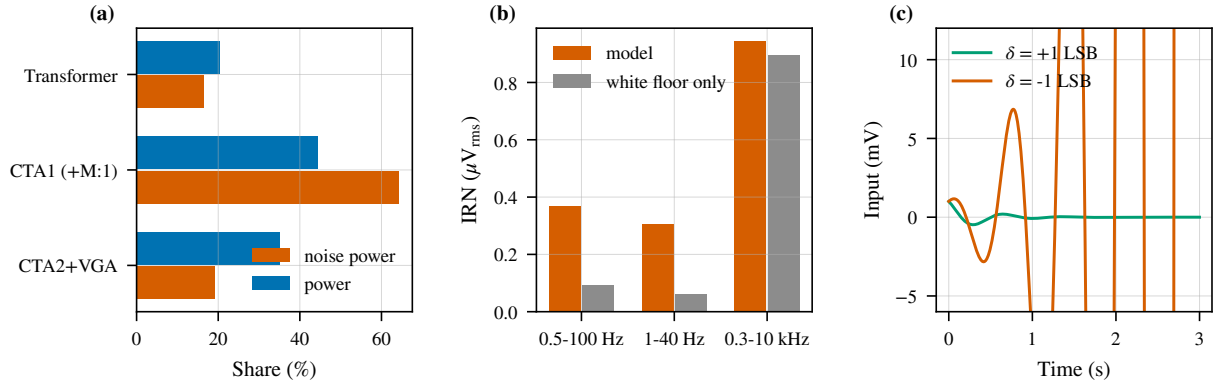


**Figure 3:** Limitations of the published design in the model. (a) Contributions to the noise power (10 kHz bandwidth) and to the power consumption, by block. (b) Integrated input-referred noise by band, with all noise sources and with the white-noise floor only. (c) Input behind a 1 nF capacitive electrode with a 1 mV initial offset, with the precharge one code below (stable) and one code above (divergent) the charge balance.

# 4 Proposed Techniques

Each technique was implemented with the same block models and simulated against the same baseline; Table 2 summarizes the results. **P4** and **P2** are the principal contributions and are complementary.

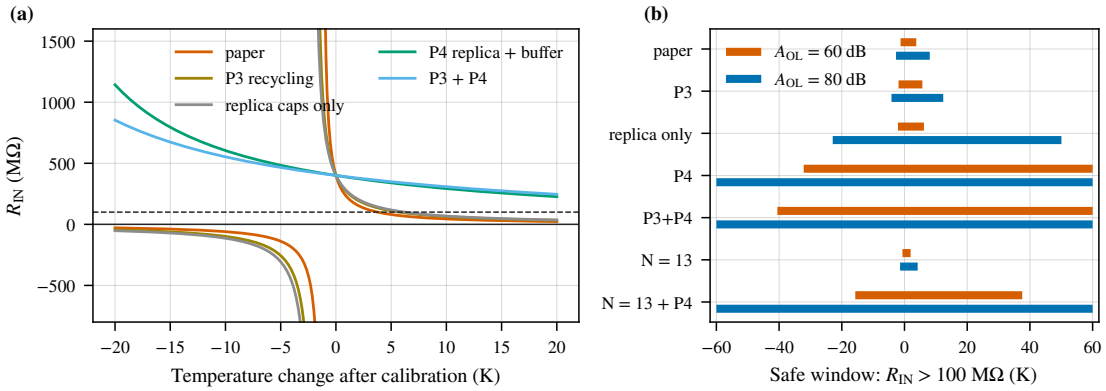


**Figure 4:** (a) $R_{IN}$ versus the temperature change after a single calibration, for the published precharge (labelled paper), **P3**, replica precharge capacitors alone, **P4**, and **P3** combined with **P4**. (b) Temperature ranges over which $R_{IN}$ remains positive and above 100 MΩ, for CTA1 open-loop gains of 60 and 80 dB, including $N = 13$ with and without **P4**.

**P4: drift-tracking precharge.** **P4** addresses both drift mechanisms. The precharge capacitors are driven by a unity-gain buffer of the transformer output instead of by CTA1 through the M:1 down-converter, which reduces the gain error from approximately $13.6/A_{OL}$ to $1/A_{OL}$; the recent SPA neural interface of [2] drives its precharge capacitors in the same way, and in the model this buffer alone reduces the drift to 22 ppm/K and widens the safe range to −5.3 to 15.6 K. The new element of **P4** is to implement the precharge capacitors with the same structure as the bottom-plate parasitic capacitance, namely replica bottom-plate capacitors with dummy switch junctions at the same bias, so that their ratio to the parasitic capacitance is maintained over temperature. With both, the drift of $\delta$ decreases to 4.0 ppm/K, and $R_{IN}$ remains above 100 MΩ from −32 K to beyond 60 K, the limit of the simulated range (Fig. 4). In the time-domain simulation at 3 K below the calibration temperature, $R_{IN}$ is 434 MΩ, compared with −290 MΩ for the published precharge, and over the 8-hour profile it remains above 377 MΩ after a single calibration. Fig. 5 shows the corresponding time-domain behaviour with a 1 nF capacitive electrode: with the published precharge the input diverges to the limit of the input range, whereas with **P4** the electrode discharges. The overhead is the bias current of the buffer, which the model does not include, and the replica capacitors must match the parasitic capacitance to approximately 0.1 %. Replica structures are common in analog design; to the author's knowledge, they have not previously been used to implement the precharge capacitors of an impedance-boosting loop, including the calibrated capacitor DAC of [2].

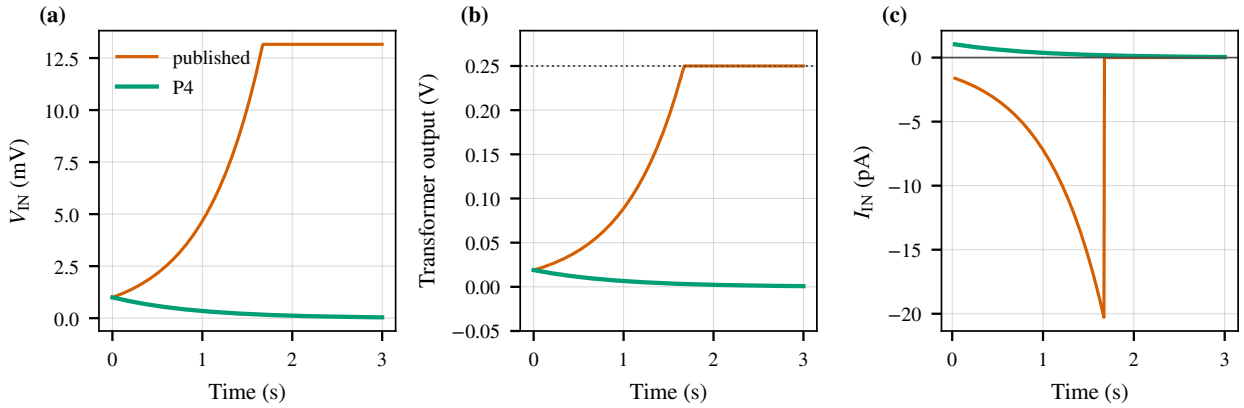


**Figure 5:** Time-domain simulation of **P4** with a 1 nF capacitive electrode and a 1 mV initial offset, 3 K below the calibration temperature, without signal or noise. (a) Input voltage. (b) Transformer output with its ±0.25 V limit (dotted). (c) Current drawn from the electrode. The model has no ICC loop, so the input node also carries the sampling capacitance, which slows both responses relative to the fabricated circuit.

**P2: chopper stabilization within the transformer.** The transformer output is chopped at 2 kHz so that CTA1 and CTA2 amplify the signal away from their 1/f noise; the signal is demodulated after CTA2,

and the ICC loop is driven from the demodulated node. The input-referred noise between 0.5 and 100 Hz decreases from 0.37 to 0.093 $\mu V_{rms}$, and between 1 and 40 Hz by a factor of 5.1 (Fig. 6(a)); time-domain simulation yields factors of 3.2 between 10 and 100 Hz and 4.2 between 4 and 40 Hz. No improvement is obtained in the spike band, where the white-noise floor dominates. The time-domain simulation also reveals an interaction with the IRB loop: when the precharge is derived from CTA1, as in [1], each chopper transition corrupts the precharge while CTA1 settles, and $R_{IN}$ collapses to 1.6 MΩ (Fig. 6(b)). Holding the precharge for 40 µs after each transition restores 386 MΩ. With a buffered precharge taken ahead of the chopper, as in **P4** and [2], no hold is required (389 MΩ). In the code-domain-multiplexed readout of [2], demodulation removes the flicker noise of all channels except the one whose Walsh code is all ones; chopping inside the SPA would also suppress it for that channel. Chopper stabilization is established in continuous-time amplifiers and was used in the discrete-time parametric amplifier of [4]; to the author's knowledge, it has not been applied to a series-parallel transformer, and the precharge-hold requirement is specific to the IRB loop.

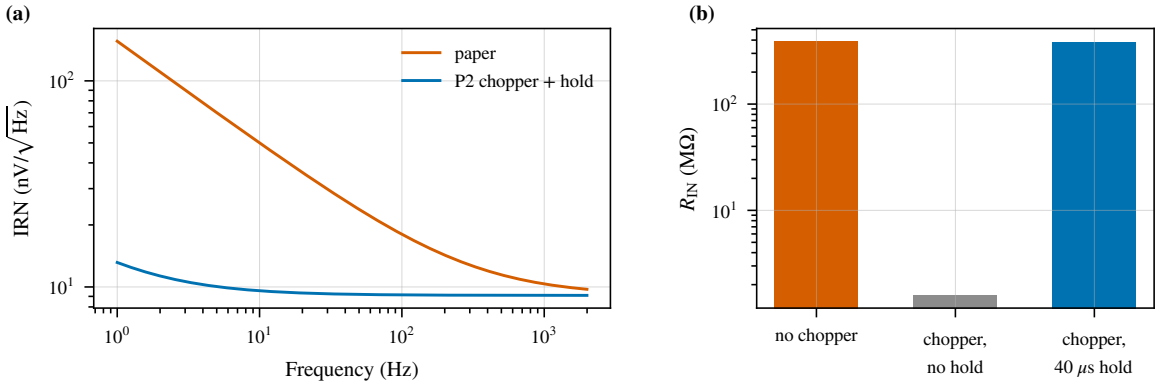


**Figure 6:** **P2**. (a) Input-referred noise density without (labelled paper) and with the chopper. (b) $R_{IN}$ at DC from the time-domain simulation with the published precharge: without chopper, with chopper and no precharge hold, and with chopper and a 40 µs hold.

**P6: delta transformer.** With a DC-coupled passive gain of 19 and a ±0.25 V transformer output swing, the input range is limited to approximately ±13 mV, which electrode offsets and stimulation artifacts of the order of 100 $mV_{pp}$ exceed. In **P6**, the bottom plates are connected during sampling to a prediction $V_P$ of the input instead of to the reference, so that the capacitors store only the prediction error. $V_P$ is subtracted at the CTA1 input and restored in the reconstruction; it is generated by a 10-bit DAC driven by an integrating servo updated every 32 cycles. The input range then equals the DAC range of ±250 mV. For a 100 $mV_{pp}$, 2 Hz artifact superimposed on a 200 µV, 40 Hz signal (Fig. 7), the published transformer saturates for 83 % of the time and only 19 % of the signal amplitude is recovered, whereas the delta transformer does not saturate and recovers the full signal. Sampling $V_P$ increases the transformer noise by 3 dB (from 4.1 to 5.8 $nV/\sqrt{Hz}$), and the DAC adds power. The technique applies the delta amplification of [5, 6] to the sampling reference of the passive stage.

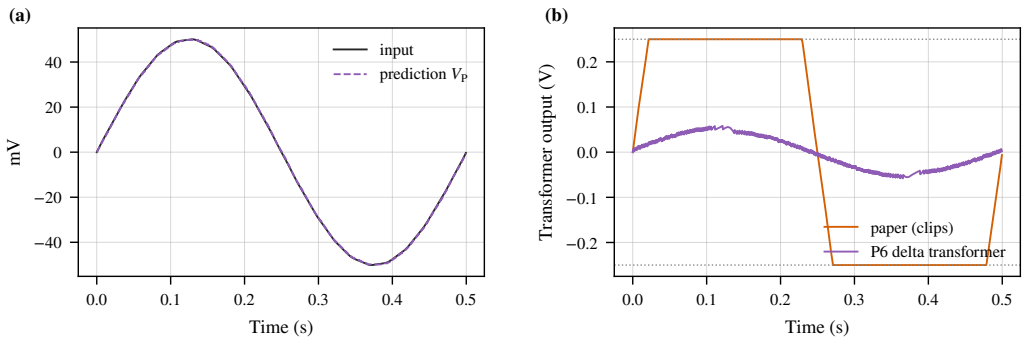


**Figure 7:** Time-domain simulation of **P6**. (a) Input (100 $mV_{pp}$ artifact and 200 µV signal) and prediction $V_P$. (b) Transformer output of the published design (labelled paper; limited to ±0.25 V, dotted) and of the delta transformer.

**P1: margin-preserving recalibration.** The calibration of [1] increments the code until the comparator decision changes, whereas the more recent interface of [2] uses the ADC output for a binary-search calibration of the precharge capacitors. With the same termination criterion, a binary search followed by a short linear search finds the code in 20 rather than 489 comparisons, reducing the time during which the electrodes are disconnected from 35 to 1.4 ms. P1 adds two elements to such a calibration: a one-code back-off, which keeps $R_{\mathrm{IN}}$ positive, and recalibration triggered by an on-chip temperature sensor. Over the 8-hour profile, $R_{\mathrm{IN}}$ then never becomes negative, compared with 12.6 % of the time without P1 (Fig. 8); this requires 59 recalibrations, and $R_{\mathrm{IN}}$ remains below 100 MΩ for 15 % of the time because the back-off shifts the safe temperature range rather than widening it. With P4, recalibration for drift is no longer necessary. The binary search itself follows [2]; the contribution of P1 is the drift budget that sets the margin and the recalibration trigger.

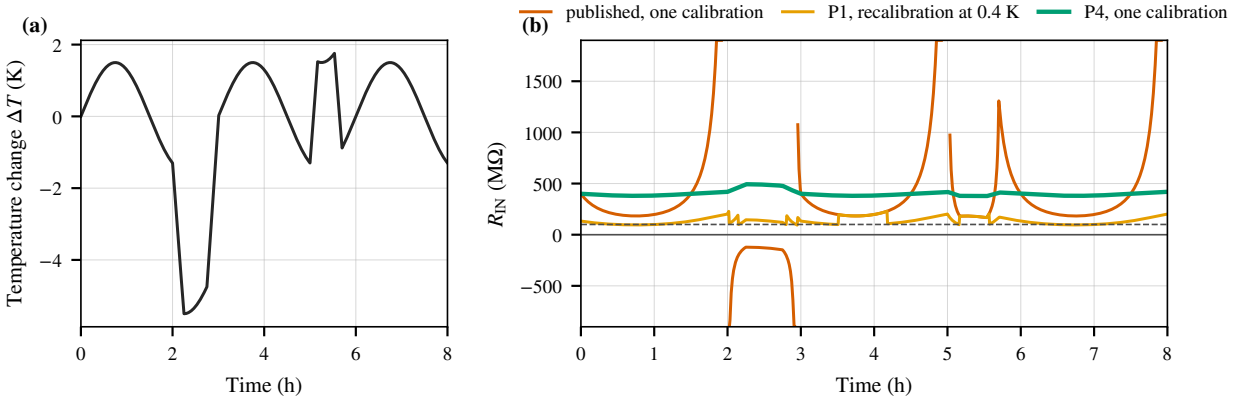


**Figure 8:** P1 and P4 over a synthetic 8-hour skin-temperature profile (closed-form model, 60 dB CTA1 open-loop gain). (a) Temperature change from the calibration temperature: a slow ±1.5 K variation, −4 K outdoors from 2 to 3 h and 3 K during exercise from 5 to 5.7 h. (b) $R_{\mathrm{IN}}$ of the published design calibrated once, with P1 (recalibration after 0.4 K, one-code back-off), and of P4 calibrated once; limited to −0.9 and 1.9 GΩ, with 100 MΩ dashed.

**P3: charge recycling between the interleaved stages.** The two transformer stages operate in alternation. Immediately before stage A is stacked, each of its bottom plates is connected to the corresponding plate of stage B, which has just left the stack at approximately $kV_{\mathrm{IN}}$; charge sharing provides half of the required parasitic charge without an active source. The precharge therefore supplies half the charge, and the drift decreases by a factor of 1.5 on its own (59 ppm/K) and to 3.3 ppm/K in combination with P4. The overhead is $2N$ switches and one short clock phase. The amplifier of [3] charged the bottom plates in several steps; here, the passive step is taken from the other interleaved stage, and the calibrated loop completes the restoration.

**P5: larger transformer ratio.** The input-referred noise of CTA1 is divided by the passive gain $2N + 1$. At constant total capacitance and clock frequency, $N = 13$ reduces the NEF from 0.31 to 0.25 and the PEF from 0.084 to 0.053 in the model (which yields 0.31 at $N = 9$, against the measured 0.34), for 9 % more power (Fig. 9). The drift budget, however, decreases approximately as $1/N^2$: with the published precharge, $N = 13$ maintains $R_{\mathrm{IN}}$ above 100 MΩ over only 2.6 K, whereas with P4 the range remains 53 K. The input range decreases by a factor of 0.7, and the ICC loop must cancel a larger capacitance. The choice of $N = 9$ in [1] follows from the trade-off between gain and input impedance (Sec. II-C); P4 relaxes this trade-off.

## 5 Limitations and Future Work

All results are obtained from behavioural simulation. The drift results scale with the open-loop gain of CTA1 and the temperature coefficients, which are not reported in [1]; with an open-loop gain of 80 dB, the safe range of the published design widens to −2.7 to 8.0 K, which remains far narrower than with P4. The model does not include the bias current of the P4 buffer, the DAC of P6 or the charge injection of the chopper switches, and the ICC loop is represented by a fit to the measured impedance rather than by a circuit model. The next step is a transistor-level implementation of P4 together with P2 in an open process

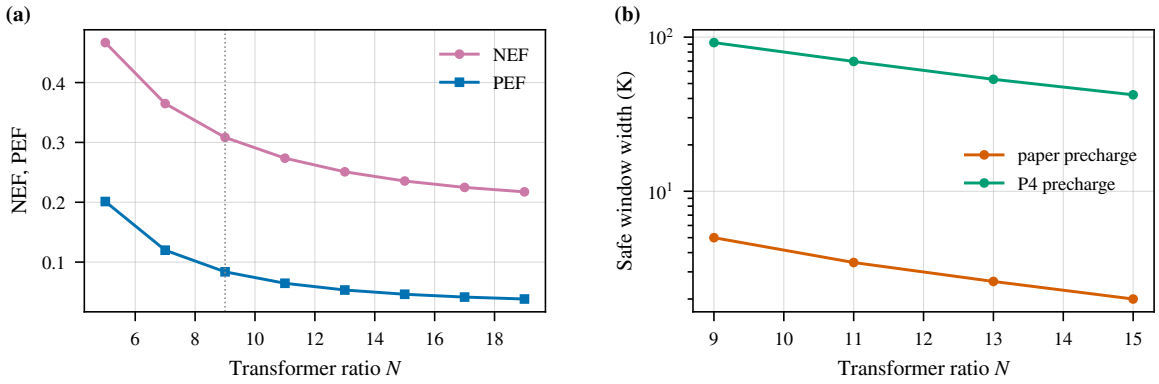


**Figure 9:** P5. (a) Model NEF and PEF versus transformer ratio $N$ (dotted line: the published $N = 9$). (b) Width of the temperature range with $R_{\mathrm{IN}} > 100\,\mathrm{M\Omega}$ versus $N$, with the published precharge (labelled paper) and with **P4**.

design kit, to verify how closely replica capacitors track the actual bottom-plate parasitic capacitance over temperature and supply voltage.

## 6 Conclusion

A behavioural model of the impedance-boosted transformer-first front end of [1], validated against its measurements, shows that the 400 MΩ input resistance is sensitive to temperature after a single calibration and that the amplifiers following the transformer dominate the low-frequency noise. A drift-tracking precharge (**P4**), which adds replica precharge capacitors to a buffered precharge such as that of [2], extends the temperature range over which the input resistance exceeds 100 MΩ from 5 K to at least 92 K, and chopper stabilization within the transformer (**P2**) reduces the noise between 0.5 and 100 Hz by a factor of four without a precharge hold. Together with the delta transformer (**P6**) and three smaller techniques, they improve the robustness, low-frequency noise and input range of the architecture while retaining its passive first gain stage.

## References

[1] G. Atzeni, C. Livanelioglu, S. Arjmandpour, and T. Jang, "An impedance-boosted transformer-first discrete-time analog front-end achieving 0.34 NEF and 389-MΩ input impedance," *IEEE J. Solid-State Circuits*, vol. 59, no. 4, 2024.

[2] Y. Chen *et al.*, "A 256-channel, 32-interconnect, 1.25-NEF, in vitro and in vivo neural interface with auto-calibrated input impedance boosting and on-chip spike detection for flexible capacitive microelectrode array," in *Proc. IEEE/JSAP Symp. VLSI Technology and Circuits*, 2026.

[3] G. Atzeni *et al.*, "An impedance-boosted switched-capacitor low-noise amplifier achieving 0.4 NEF," in *Proc. IEEE Symp. VLSI Technology and Circuits*, 2022.

[4] T. Jang *et al.*, "A noise-efficient neural recording amplifier using discrete-time parametric amplification," *IEEE Solid-State Circuits Lett.*, vol. 1, no. 11, 2018.

[5] K. Jeong, C. Livanelioglu, J. Liao, I. Lee, and T. Jang, "A 3.47 NEF 175.2dB FOMs direct digitization front-end featuring delta amplification for enhanced dynamic range and energy efficiency in bio-signal acquisition," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2025.

[6] K. Jeong, C. Livanelioglu, J. Liao, I. Lee, and T. Jang, "A 3.47 NEF 175.2-dB FoMs direct digitization front-end featuring delta amplification noise-shaping SAR ADC for biosignal acquisition," *IEEE J. Solid-State Circuits*, vol. 61, no. 9, 2026.