Modelling and Improving an Impedance-Boosted Transformer-First Discrete-Time Analog Front End

Nithin Purushothama*
nithinpurushothama@gmail.com • nithinpuru.com

November 2026

Abstract

The transformer-first analog front end (AFE) of Atzeni *et al.* [1] uses a discrete-time series-parallel switched-capacitor transformer as a passive, nearly flicker-free gain stage and recovers the input impedance it costs with two active loops, reaching 0.34 NEF, 0.1 PEF and 389 MΩ at 1 kHz. This report rebuilds that front end as connected behavioural blocks (a charge-level model of the transformer with its bottom-plate parasitics, the precharge path, the continuous-time amplifiers, the cancelling loop and the calibration), integrates them into the full chain, and checks the result against the published measurements: transformer gain 25.58 dB (paper 25.6 dB), NEF/PEF 0.340/0.102 with the measured noise, and a simulated kT/C floor of 3.79 nV/$\sqrt{\text{Hz}}$ (paper 4 nV/$\sqrt{\text{Hz}}$). The integrated model then shows where the design is limited: CTA1 makes 64 % of the noise power, the CTAs' $1/f$ noise sets the noise below about 300 Hz, and the boosted input resistance rests on restoring the parasitic charge to within about 0.01 %, so with the published calibration rule it stays positive and above 100 MΩ only within a few kelvin of the calibration temperature. Six improvements are modelled with the same blocks: a noise-robust, margin-keeping recalibration; a drift-tracking precharge (22x less drift, safe window −32 K to 60 K); a chopper merged into the transformer (3.2 to 4.2x less noise below 100 Hz, provided the precharge is held across chopper transitions); charge recycling between the interleaved stages; a larger transformer ratio paid for by the robust precharge (NEF 0.31 → 0.25); and a delta transformer that tolerates a 100 mV_{pp} artifact without clipping, at a 3 dB cost in the transformer's noise. All results are simulations of a behavioural model whose unpublished parameters are stated as assumptions.

1 Introduction

Biopotential front ends spend most of their power in the first amplifier, whose noise per unit current is bounded by the transconductance efficiency of its transistors; continuous-time designs therefore stay above an NEF of one [2, 1]. Atzeni *et al.* avoid that bound with a passive first stage: N capacitors charged in parallel and restacked in series give a voltage gain of $N + 1$ for the cost of switching them [3, 1]. Their JSSC 2024 front end adds an input-resistance-boosting (IRB) loop, which precharges the transformer's bottom-plate parasitics from the next amplifier, and an input-capacitance-cancelling (ICC) loop, and reaches the lowest reported NEF and PEF [1].

The aim here is to understand that design well enough to improve it. The contributions are:

- a block-level behavioural model of the front end, from a charge-level transformer to the calibration, integrated into the full chain and checked against the paper's measurements (Secs. 3 and 4);
- a bottleneck analysis with the integrated model (Sec. 5);
- six proposed improvements, each implemented in the same blocks and evaluated against the same baseline (Sec. 6).

Nothing here was fabricated or measured; quantities the paper does not give are assumptions, listed in Table 1.

*Interactive version of the model: `https://nithinpuru.com/spa_afe/`.

2 The front end

Fig. 1 shows the AFE as modelled. In the sample phase Φ_1 each capacitor $C_{S,k}$ holds V_{IN}; in Φ_2 the capacitors are stacked on the input, giving $(N+1)V_{\mathrm{IN}}$ per path and $(2N+1)V_{\mathrm{DM}}$ differentially ($N = 9$, 25.6 dB measured). The input-referred kT/C noise of the transformer and its NEF are [1, 4]

$$v^2_{\mathrm{rms,in}} = \frac{N^2}{A^2}\frac{2kT}{C_{\mathrm{tot}} f_S}\frac{\pi}{2} f_{\mathrm{BW}}, \qquad \mathrm{NEF}_{\mathrm{SPA}} = \frac{N}{A}\sqrt{\frac{C_{\mathrm{SW}} V_{\mathrm{SW}}}{2U_T C_{\mathrm{tot}}}}, \tag{1}$$

so the noise efficiency depends on a capacitance ratio rather than on g_m/I_D. The cost is the input: the bottom-plate parasitic C_{BP} of capacitor k has to charge to about kV_{IN} every cycle, which takes charge from the sampling capacitors and, in turn, from the input, an input resistance $R_{\mathrm{IN}} \propto 1/(C_{\mathrm{BP}} f_S)$ [1, eq. (15)]. The IRB loop shares each parasitic with a precharge capacitor charged to βV_{OUT} ($\beta = A_{\mathrm{CTA1}}/M$) just before stacking, sized as $C_{\mathrm{PC},i} = iC_{\mathrm{BP}}/((N+1)\beta - i)$ [1, eq. (21)]; only the total restored charge matters, so one capacitor, C_{PCN}, is trimmed by a foreground search that floats the input and steps the code by one LSB until the input voltage stops drooping [1, Sec. III-E]. The ICC loop feeds C_{PF} from CTA2 back to the input so that $(A_{\mathrm{tot}} - 1)C_{\mathrm{PF}}$ cancels the input capacitance, leaving about 200 fF for stability [1, Sec. III-D].

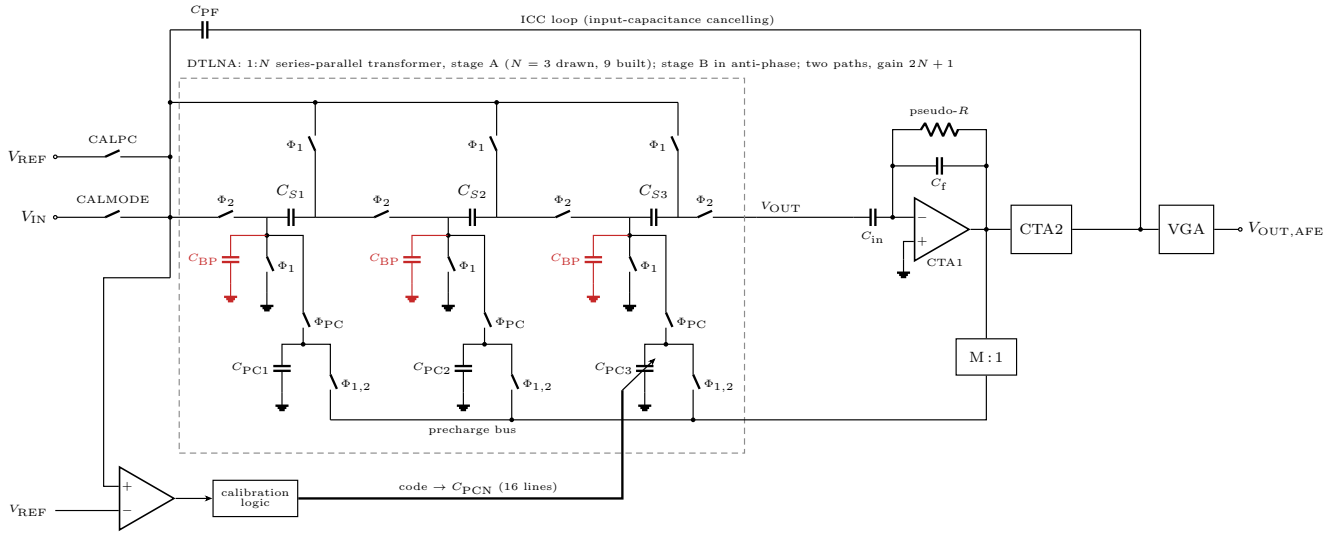

Figure 1: The front end of [1] as modelled, drawn after its Figs. 4 and 5(a): one 1:N series-parallel stage ($N = 3$ drawn), bottom-plate parasitics (red), precharge capacitors and bus (IRB loop), CTA1 with capacitive feedback, M:1 downconverter, CTA2, VGA, the ICC loop and the calibration. The proposed changes are drawn on the same circuit in Fig. 4.

3 Block models and integration

Parameters. Table 1 lists what the model takes from the paper and what it assumes. Two inferences anchor the rest. First, the measured $4\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$ floor fixes $C_S f_S = 2.6 \times 10^{-5}\,\mathrm{F/s}$ through (1); with the $6\,\mathrm{fF}/\mu\mathrm{m}^2$ of the paper's Fig. 3, the $0.15\,\mathrm{mm}^2$ transformer holds 36 capacitors of about 26 pF, so $f_S \approx 1\,\mathrm{MHz}$, which with the inferred 75 nW gives $\mathrm{NEF}_{\mathrm{SPA}} \approx 0.1$, the value the paper states. Second, the $160\,\mathrm{k}\Omega$ measured "before calibration" is read as precharge on with C_{PCN} at code 0, which gives $C_{\mathrm{BP}} = 39.6\,\mathrm{fF}$ ($\alpha = 0.15\,\%$) and an unboosted $R_0 = 45.6\,\mathrm{k}\Omega$.

Transformer. One stage is solved at the charge level: in Φ_1 the capacitors connect to the input and the reference; in Φ_2 charge conservation on every floating island (top plate of capacitor k, bottom plate and C_{BP} of capacitor $k+1$, and the load) gives the stack voltages, and the charge drawn from the input follows from the island that holds the input. With $C_{\mathrm{BP}} = 0$ the gain is exactly $N+1$ and no input charge flows; with the precharge of [1, eq. (21)] the gain returns exactly to $N+1$ and the input charge to zero, which confirms that a single trimmed capacitor can null the loss. The charge each bottom plate costs the input is

Table 1: Model parameters: from the paper, read from its figures, derived, or assumed.

Quantity	Value	Source
N, gains, bandwidth, supplies	9; DTLNA/CTA1/CTA2 25.6/22/7 dB; 10 kHz; 0.6/1 V	paper, Sec. IV, Table I
Noise floors, IRN	4/8.6/9.1 nV/$\sqrt{\text{Hz}}$; 1.36 μV_{rms}	paper, Sec. IV
$1/f$ corners (transformer/CTA1/CTA2+VGA)	6/131/1865 Hz	fitted to Fig. 10(c) at 10 Hz
Power split	75/165/130 nW	derived from Table I NEF, PEF, power
f_S, C_S	1 MHz, 25.8 pF	derived (floor and area)
C_{BP}, R_0	39.6 fF, 45.6 kΩ	derived (160 kΩ at code 0)
M, C_L, code width, C_{LSB}	8, 1 pF, 16 bits, 0.127 fF	assumed; LSB set so calibration lands at 400 MΩ
CTA1 A_{OL}, temperature coefficients	60 dB (also 80 dB); A_{OL} −0.4 %/K, C_{BP} 100 ppm/K, MOM 30 ppm/K	assumed
ICC and precharge-path terms	residual 200 fF; C_c, loop poles, high-pass	paper (residual); fitted to Fig. 10(b)

weighted by about k (it passes through capacitors $1 \dots k-1$ and is refilled in the next sample phase), so the model draws about 1.4 times the input charge of [1, eq. (15)] for the same α; because R_0 is fitted to the measurement, this changes only the implied α.

IRB, calibration, CTAs, ICC. With δ the fraction of the parasitic charge that the precharge does not restore, $R_{IN} = R_0/\delta$. δ follows from eq. (20) of [1] for every plate, the CTA1 closed-loop gain (finite A_{OL}), the M:1 ratio and the capacitor temperature coefficients. The calibration floats the input on the residual capacitance, lets it droop for 64 cycles and compares it with V_{REF} (offset 1 mV, noise 0.5 mV). The CTAs have capacitive feedback, one pole each and white plus $1/f$ noise. The input admittance combines the IRB conductance with a fitted precharge-path high-pass, the residual capacitance and the ICC loop with two poles.

Integration. The time-domain model runs both paths and both interleaved stages of the charge-level transformer at 1 MHz with kT/C noise, CTA1 (discretised exactly for the held transformer output), the M:1 precharge from CTA1's output, CTA2 and the VGA. The frequency-domain model gives the gain, noise budget, IRN, NEF/PEF and input impedance. A browser port of the frequency-domain parts is checked against 72 values from the Python model.

4 Validation

Table 2 and Fig. 2 compare the integrated model with the paper. The time-domain loop and the block equations agree within 2.5 % on the input resistance at every code. The integrated IRN is 9 % below the measured 1.36 μV_{rms}; the paper does not state its integration limits. The input-impedance fit has 0.06 decade rms error.

Table 2: Integrated model against the paper (time domain where stated).

Quantity	Model	Paper
Transformer gain	25.58 dB; time domain 18.9999 V/V	25.6 dB
R_{IN} before / after calibration	160 kΩ / 400 MΩ; time domain 158 kΩ / 390 MΩ	160 kΩ / about 400 MΩ
Transformer kT/C floor	3.79 nV/$\sqrt{\text{Hz}}$ (time domain)	4 nV/$\sqrt{\text{Hz}}$
IRN, 1 Hz up, single-pole 10 kHz	1.23 μV_{rms}	1.36 μV_{rms}
NEF / PEF (measured IRN)	0.340 / 0.102; chain 0.270	0.34 / 0.1; 0.27
$\lvert Z_{IN}\rvert$ at 1 kHz	453 MΩ (fit)	389 MΩ
Calibration	sign flip between adjacent codes; 489 comparisons	linear search

Three open points came out of the validation. (i) The noise floor and the area imply about 230 pF

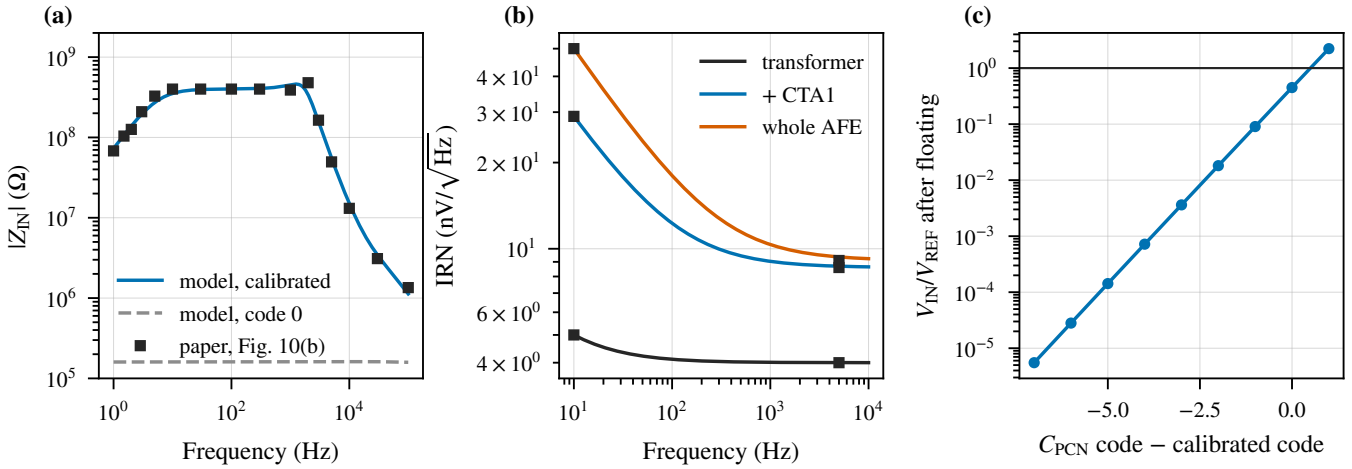

Figure 2: Validation. (a) $|Z_{IN}|$: model calibrated and at code 0 (lines) and the paper's Fig. 10(b) (squares). (b) Input-referred noise of the transformer, transformer + CTA1 and the whole AFE (lines); the paper's values at 10 Hz (read from Fig. 10(c)) and its stated floors (squares). (c) Calibration: input voltage relative to V_{REF} after the floating interval against the code; the flip marks the calibrated code.

of sampling capacitance per stage, yet $|Z_{IN}|$ before calibration stays flat to 100 kHz, which bounds the uncancelled input capacitance to a few picofarads. (ii) The impedance fit rings with a high source resistance (pole $Q \approx 1.1$ at 10 MΩ) whereas the paper's circuit simulation is critically damped up to 34 MΩ; the fit is therefore used for magnitudes only. (iii) What "before calibration" means in Fig. 10(b) sets R_0; the alternative reading (precharge off) gives $R_0 = 160$ kΩ and 3.5 times more drift margin than reported below.

5 Bottlenecks

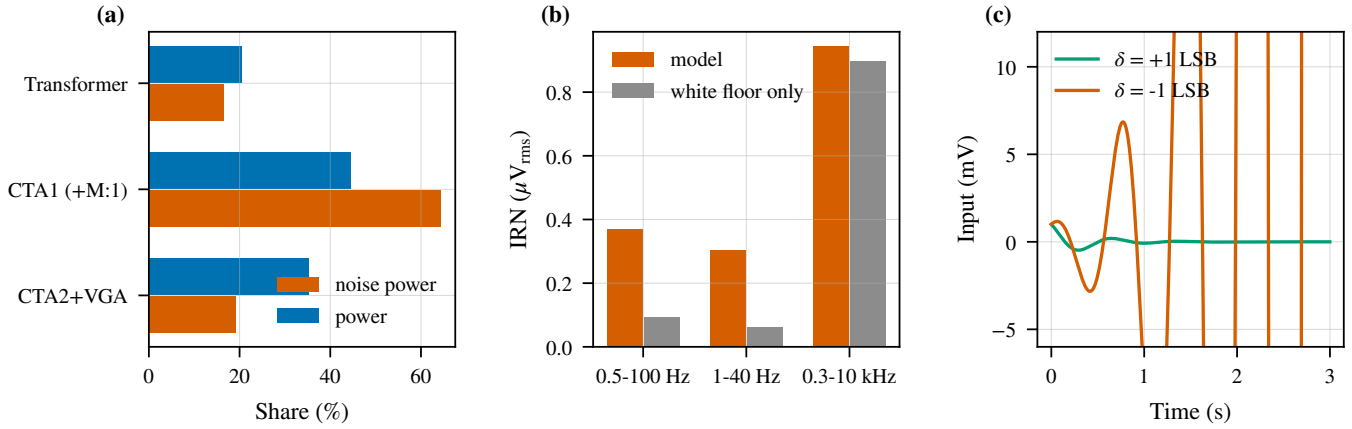

Figure 3: Bottlenecks of the baseline. (a) Shares of noise power (10 kHz band) and of power by block. (b) Integrated IRN by band, model and white floor only. (c) Input behind a 1 nF capacitive electrode with the precharge one LSB under (stable) and one LSB over (runs away) the balance, 1 mV initial offset.

Noise and power. CTA1 makes 64 % of the noise power in the 10 kHz band and takes 45 % of the power; CTA2 and the VGA take 35 % of the power for 19 % of the noise; the transformer makes 17 % of the noise for 20 % of the power (Fig. 3a). Below about 300 Hz the CTAs' $1/f$ noise dominates: from 0.5 to 100 Hz the model integrates to 0.37 μV_{rms} against a white floor of 0.09 μV_{rms} (Fig. 3b).

Input-resistance drift. One calibration step changes δ by 2.3×10^{-4}. With CTA1's closed-loop gain error $(1 + A_{CTA1})/A_{OL}$ and the C_{BP}/MOM temperature-coefficient mismatch, δ drifts by 9.0×10^{-5} K at $A_{OL} = 60$ dB (60 % from the gain error, 40 % from the capacitors) and 4.2×10^{-5} K at 80 dB. With the published rule (stop at the last positive code), R_{IN} stays positive and above 100 MΩ only from −1.2 K to 3.8 K around the calibration temperature (−2.7 K to 8.0 K at 80 dB); cooling beyond that over-restores

the charge and makes R_{IN} negative. With a capacitive (DC-isolated) electrode, the AFE's own input conductance sets the input's DC level, and one LSB of over-restoration drives the input to the rail within about 3 s (Fig. 3c). This matters for the capacitive arrays this group now uses [5, 6].

6 Proposed improvements

Table 3: Published design against the proposal (summary of this section). Proposed values at $N = 9$ except the **P5** rows; drift numbers assume a 60 dB CTA1 loop gain and the temperature coefficients of Table 1.

	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 $R_{\mathrm{IN}} > 100$ 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 δ	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 (electrodes disconnected)	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 $R_{\mathrm{IN}} > 100$ 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/$\sqrt{\mathrm{Hz}}$ (Sec. IV)	4.1 nV/$\sqrt{\mathrm{Hz}}$	5.8 nV/$\sqrt{\mathrm{Hz}}$	+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**

Each proposal answers one bottleneck and was implemented in the same blocks; Table 4 summarises the modelled results.

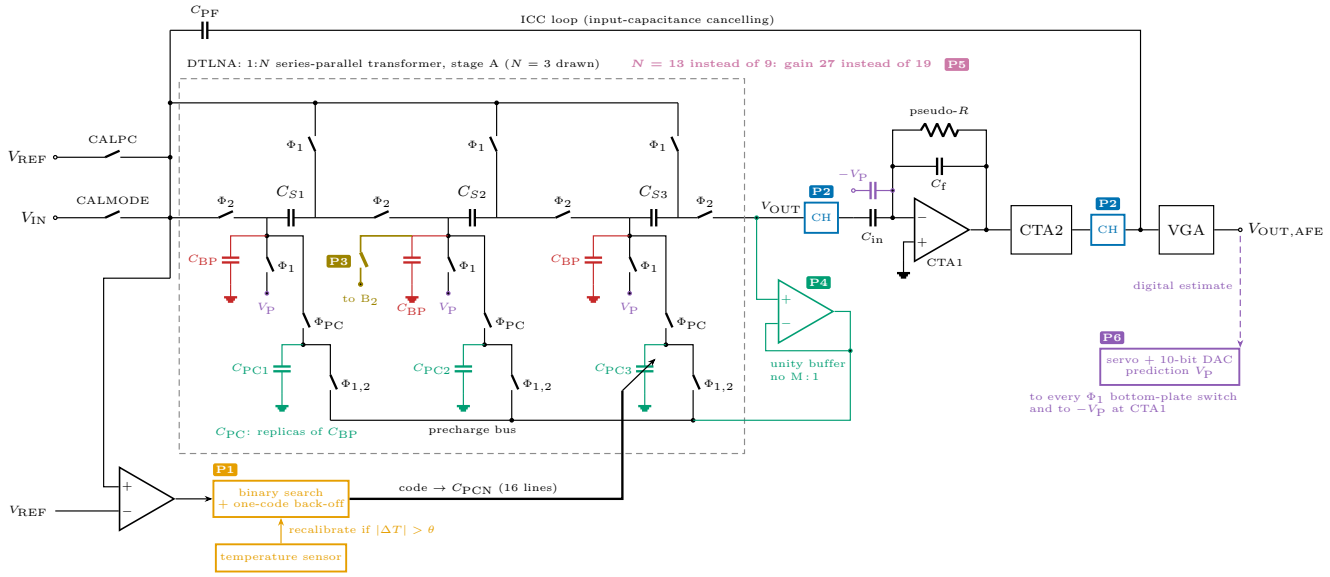

Figure 4: The proposed front end: the published circuit (black, parasitics red) with the six changes in colour. **P4** takes the precharge from a unity buffer of the transformer output before the **P2** chopper, so the chopper needs no precharge hold.

Table 4: Proposed improvements, modelled with the same blocks (CTA1 $A_{OL} = 60\,\text{dB}$ unless stated).

Proposal	Result	Baseline
P1 margin-keeping recalibration (Sec. 6.1)	8 h profile: never negative; 59 recalibrations, 85 ms disconnected; $R_{IN} < 100\,\text{M}\Omega$ 15 % of the time	calibrated once: negative 12.6 % of the time
P2 chopper in the transformer (Sec. 6.2)	IRN 10 Hz to 100 Hz $0.068\,\mu\text{V}_{\text{rms}}$, 4 Hz to 40 Hz 0.031; R_{IN} 386 MΩ	0.218 and $0.131\,\mu\text{V}_{\text{rms}}$; 390 MΩ
P3 charge recycling (Sec. 6.3)	drift 5.9×10^{-5} K; half the precharge charge from CTA1	9.0×10^{-5} K
P4 drift-tracking precharge (Sec. 6.4)	drift 4.0×10^{-6} K; window −32 K to 60 K; 8 h minimum 377 MΩ with one calibration	−1.2 K to 3.8 K
P5 larger ratio, $N = 13$ (Sec. 6.5)	NEF 0.251, PEF 0.053, 403 nW; with **P4** window −16 K to 38 K	0.308, 0.084, 370 nW
P6 delta transformer (Sec. 6.6)	$100\,\text{mV}_{\text{pp}}$ artifact: no clipping, 200 µV signal recovered; floor $5.8\,\text{nV}/\sqrt{\text{Hz}}$	clips 83 % of the time; $4.1\,\text{nV}/\sqrt{\text{Hz}}$

6.1 P1: margin-keeping, temperature-triggered recalibration

Motivation. The calibration lands anywhere between 0 and 1 LSB from the sign change, and a few kelvin of drift cross it (Sec. 5).

Method. Find the code by binary search followed by a 3-step linear search that ends with the published criterion; then step back one code, which guarantees one LSB of margin; recalibrate when an on-chip temperature sensor reports a change above a threshold since the last calibration. A noisy-comparator variant would use probabilistic bisection [7].

Result. The search needs 20 comparisons instead of 541 on average (1,215 worst case over 100 process draws of ±20 % in C_{BP}), with the same code within one LSB. Over an 8-hour synthetic skin-temperature profile (Fig. 5: ±1.5 K slow swing, −4 K outdoors, 3 K during exercise), calibrating once leaves R_{IN} negative 12.6 % of the time. Recalibrating at 0.4 K with the back-off keeps it positive throughout, with 59 recalibrations and 85 ms of electrode-disconnected time (about 2.1 s with the linear search), but the coarse LSB leaves R_{IN} between 100 and 200 MΩ after each calibration and below 100 MΩ 15 % of the time. **P1** makes the

published design safe; it does not by itself keep R_{IN} high.

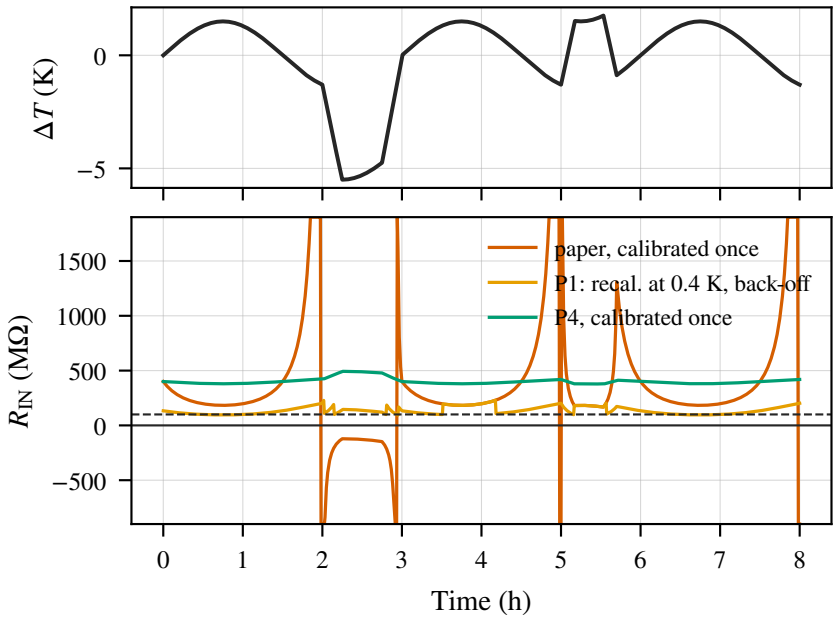

Figure 5: P1: temperature change over an 8-hour profile (top) and R_{IN} for the published design calibrated once, **P1** (recalibration at 0.4 K with one-code back-off) and **P4** calibrated once (bottom; clipped at −0.9 and 1.9 GΩ).

6.2 P2: chopper merged into the transformer, with precharge hold

Motivation. The transformer is nearly flicker-free; CTA1 and CTA2 are not, and their $1/f$ noise sets the noise in EEG/ECG bands (Fig. 3b).

Circuit. Alternate the polarity of the transformer output (swap the two paths' outputs) at $f_{\mathrm{ch}} = 2\,\mathrm{kHz}$, amplify with CTA1 and CTA2, and demodulate after CTA2 [8]. The precharge path must see the un-chopped polarity, so the CTA1 output is demodulated before the M:1, and the ICC feedback is taken after the demodulator.

Design iteration. The first version demodulated the precharge path directly. In the time-domain loop, CTA1 cannot follow the polarity flips, the precharge is wrong for several microseconds after every flip and R_{IN} collapsed from 390 MΩ to 1.6 MΩ. Holding the precharge for 40 µs after each flip (the M:1 skips its update) restores R_{IN} to 386 MΩ (Fig. 6b).

Result. In the time-domain loop, the IRN falls from 0.218 to 0.068 $\mu\mathrm{V_{rms}}$ in 10 Hz to 100 Hz (3.2 times) and from 0.131 to 0.031 $\mu\mathrm{V_{rms}}$ in 4 Hz to 40 Hz (4.2 times), for a 1.1 dB gain loss from CTA2's bandwidth. Above 300 Hz the white floor dominates and the benefit vanishes, so **P2** suits low-frequency applications; the hold adds a small lag that the ICC loop has to absorb.

6.3 P3: passive charge recycling between the interleaved stages

Motivation. Every cycle, the stage that leaves the stack dumps its bottom-plate charge to ground while the other stage's IRB supplies the same charge again.

Circuit. Just before stage A stacks, connect each of its bottom plates to the matching bottom plate of stage B, which has just left the stack at about kV_{IN}. The two equal parasitics share charge and A's plates start at about half their target. The active IRB supplies only the other half, with precharge capacitors resized accordingly.

Result. Implemented in the time-domain loop (2N extra switches, one short phase), the drift falls from 9.0×10^{-5} to 5.9×10^{-5} K and CTA1 supplies half the precharge charge. The safe window grows from −1.2 K to 3.8 K to −1.9 K to 5.7 K (Fig. 7). The idea differs from the multiphase stepwise charging of [9] in

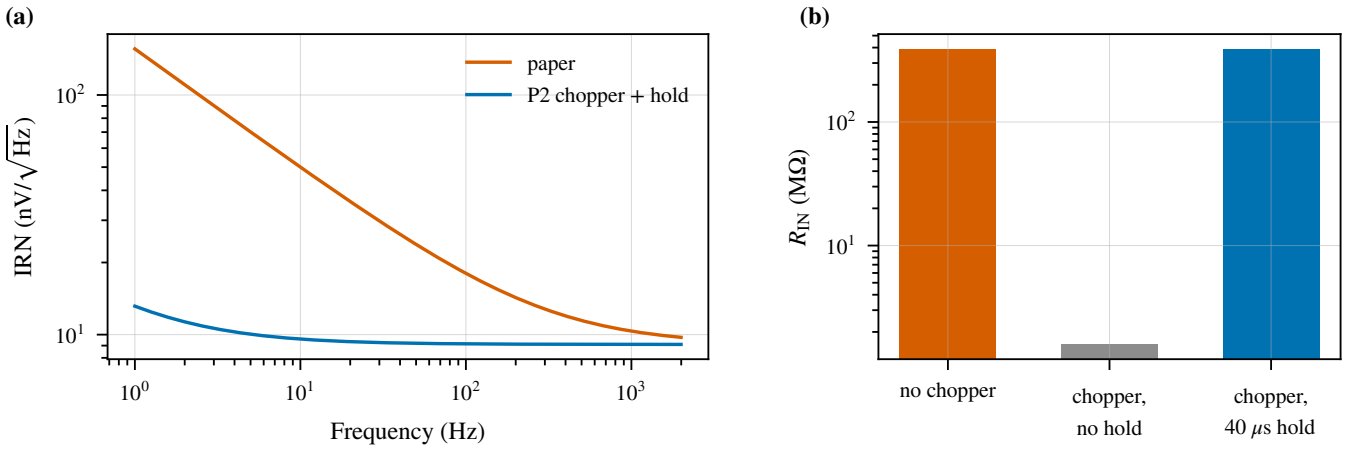

Figure 6: P2. (a) Input-referred noise without and with the chopper (model). (b) R_{IN} at DC from the time-domain loop: no chopper, chopper without precharge hold, chopper with a 40 µs hold.

that the passive step borrows charge from the other interleaved stage and the calibrated IRB finishes the job; more steps through small reservoirs would recycle more.

6.4 P4: drift-tracking precharge

Motivation. The two drift terms are CTA1's closed-loop gain error, which enters β one to one, and the temperature-coefficient mismatch between the MOM precharge capacitors and the parasitic.

Circuit. Build the precharge capacitors as replicas of the bottom-plate parasitic (the same bottom-plate-to-substrate structure plus dummy switch junctions at the same bias), so the ratio C_{PC}/C_{BP} holds over temperature; and drive them from a unity-feedback buffer of the transformer output (gain error $1/A_{OL}$, no M:1) instead of from CTA1, whose 22-dB closed-loop gain carries an error of about $13.6/A_{OL}$.

Result. Implemented in both the block and the time-domain models (agreement within 3 % at 2 K), the drift falls to 4.0×10^{-6} K (22 times less), the safe window becomes −32 K to 60 K (80 dB: the full ±60 K range), and over the 8-hour profile one calibration keeps $R_{IN} \geq 377$ MΩ (Figs. 5 and 7). Replica capacitors alone help 1.7 times; replicas plus an 80 dB CTA1 help 17 times, an alternative if the buffer's power is too high. The buffer is the main cost: it must charge the precharge capacitors every cycle and will cost bias current that the model does not count.

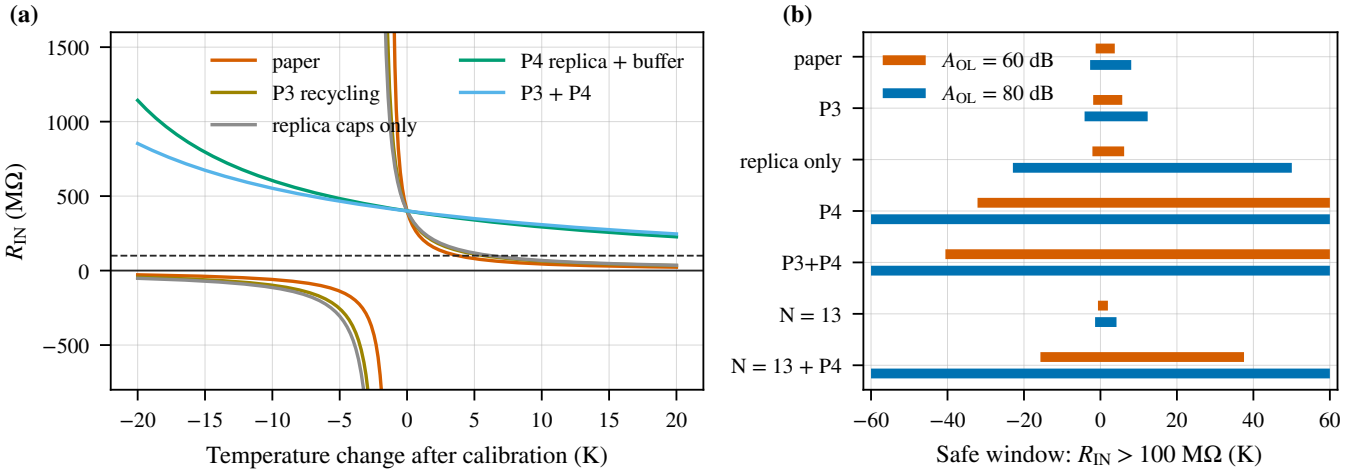

Figure 7: P3 and **P4**. (a) R_{IN} against temperature change after calibration ($A_{OL} = 60$ dB). (b) Safe windows ($R_{IN} > 100$ MΩ, positive) for all variants at $A_{OL} = 60$ and 80 dB, including $N = 13$ with and without **P4**.

6.5 P5: spend the robustness on a larger transformer ratio

Motivation. CTA1's noise is divided by the passive gain $2N+1$, but the bottom-plate charge grows about as N^3, so for a fixed R_{IN} the drift budget shrinks about as $1/N^2$.

Result. At the same total sampling capacitance and clock, $N = 13$ gives NEF 0.251 and PEF 0.053 (from 0.308 and 0.084 at $N = 9$) for 9 % more power; $N = 15$ gives 0.235 and 0.046 (Fig. 8a). The cost is the drift budget: with the published precharge the safe window at $N = 13$ shrinks to −0.6 K to 2.0 K; with **P4** it is −16 K to 38 K (Fig. 8b). Implemented in the time-domain loop, $N = 13$ gives a gain of exactly 27 and matches the block model within 5 %. The input range shrinks to 70 %, and $(N+1)^2 C_L$ adds capacitance for the ICC loop to cancel.

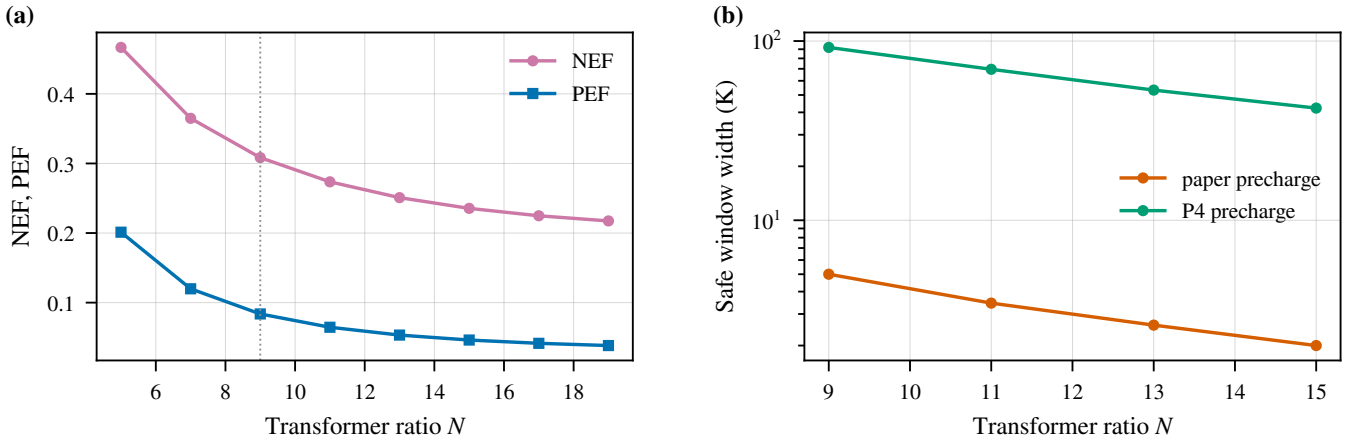

Figure 8: P5. (a) NEF and PEF against transformer ratio N (model; dotted line: published $N = 9$). (b) Width of the safe window against N with the published precharge and with **P4**.

6.6 P6: delta transformer

Motivation. The DC-coupled gain of 19 limits the input range to roughly ±13 mV for a ±0.25 V transformer output; electrode offsets and stimulation artifacts (about 100 $\mathrm{mV_{pp}}$ for transcranial current stimulation [10]) saturate it. The group's delta-amplification front end [10, 11] addresses dynamic range in a direct-digitisation ADC; this proposal applies the prediction idea to the passive transformer (the JSSC 2026 paper was not available to the author, so the overlap is unknown).

Circuit. In Φ_1 the bottom plates go to a prediction V_{P} instead of the reference, so each capacitor holds the prediction error $e = V_{\mathrm{IN}} - V_{\mathrm{P}}$; the stack still sits on the input, so the transformer output is $V_{\mathrm{IN}} + 2Ne$, and subtracting V_{P} at the CTA1 input leaves $(2N+1)e$. A 10-bit DAC over ±0.25 V sets V_{P}, updated every 32 cycles by an integrating servo on the reconstructed input $V_{\mathrm{P}} + e$. Because the parasitics now swing by ke instead of kV_{IN}, the IRB only has to restore charge for the error, and the precharge (referenced to V_{P}) comes from CTA1 as before.

Result. With a 50 mV 2 Hz artifact plus a 200 µV 40 Hz signal, the published transformer clips 83 % of the time and the signal is lost; the delta transformer never clips (output within ±59 mV) and recovers both the signal (200.0 µV) and the artifact (50.0 mV) (Fig. 9). The cost is noise: the prediction node's sampled noise is common to the N capacitors of a path and adds about as much as the capacitors themselves, raising the transformer floor from 4.1 to 5.8 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ (+3 dB, about 9 % on the whole AFE). The DAC must settle the bottom-plate charge each cycle, which costs power the model does not count.

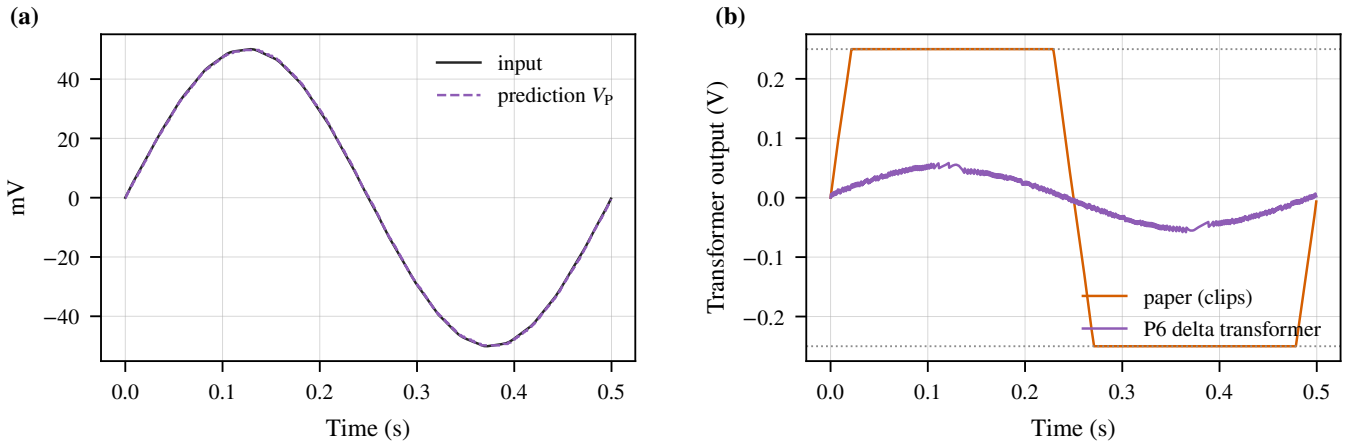

Figure 9: P6. (a) Input (artifact + signal) and the prediction V_P. (b) Transformer output of the published design (clipped at ±0.25 V, dotted) and of the delta transformer.

7 Discussion and limitations

The two strongest proposals are **P4**, which removes the main reason the boosted input resistance is fragile, and **P2**, which removes the main noise source of low-frequency recordings; **P5** shows how the first can be traded for noise efficiency. **P1** is cheap and makes the published design safe with capacitive electrodes; **P3** is modest. **P6** extends the input range by an order of magnitude at a 3 dB transformer-noise cost.

All numbers are model results. The most important assumptions are CTA1's open-loop gain and the temperature coefficients (they scale every drift number), the meaning of the measured 160 kΩ, and the bias currents of the buffer (**P4**) and DAC (**P6**), which the model does not count. The ICC model is a fit to measured magnitudes, not a circuit model, and its damping disagrees with the paper's simulation. The next step is a transistor-level check of the precharge path (temperature coefficients of C_{BP} and of a replica in an open PDK such as SKY130 or IHP SG13G2), followed by the chopper hold and the delta transformer's reference settling. The group's recent auto-calibrated boosting [6] and common-mode cancelling transformer [12] were known to the author only from their abstracts.

8 Conclusion

A block-level model of the transformer-first AFE of [1], integrated and checked against its measurements, shows that its noise efficiency is now limited by the amplifiers after the transformer, and that its boosted input resistance rests on a charge balance of about 0.01 % that drifts with temperature. Six improvements, modelled with the same blocks, address these limits; the drift-tracking precharge and the chopper merged into the transformer give the largest modelled gains.

Acknowledgments. The original circuit and all measured data are due to G. Atzeni, C. Livanelioglu, S. Arjmandpour and T. Jang [1]; any errors in the model are the author's. The model, figures and this report were developed with the help of an AI coding assistant (Claude, Anthropic), which wrote and ran the simulation scripts under the author's direction.

A Reproducibility

The Python model (`tfafe`: transformer, blocks, system, proposed), its 20 tests and the scripts for every number and figure here are in the author's repository; `python run_all.py` (25 s) and `python make_figs.py` (about 100 s) regenerate them. The browser version is a port checked against 72 Python values.

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] M. S. J. Steyaert and W. M. C. Sansen, "A micropower low-noise monolithic instrumentation amplifier for medical purposes," *IEEE J. Solid-State Circuits*, vol. 22, no. 6, 1987.

[3] G. Atzeni, A. Novello, G. Cristiano, J. Liao, and T. Jang, "A 0.45/0.2-NEF/PEF 12-nV/$\sqrt{\text{Hz}}$ highly configurable discrete-time low-noise amplifier," *IEEE Solid-State Circuits Lett.*, vol. 3, 2020.

[4] B. Murmann, "Thermal noise in track-and-hold circuits: Analysis and simulation techniques," *IEEE Solid-State Circuits Mag.*, vol. 4, no. 2, 2012.

[5] W. Choi, Y. Chen, D. Kim, S. Weaver, others, and T. Jang, "A 1,024-channel, 64-interconnect, capacitive neural interface using a cross-coupled microelectrode array and 2-dimensional code-division multiplexing," in *Proc. IEEE Symp. VLSI Technology and Circuits*, 2023.

[6] 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.

[7] R. Waeber, P. I. Frazier, and S. G. Henderson, "Bisection search with noisy responses," *SIAM J. Control Optim.*, vol. 51, no. 3, 2013.

[8] C. C. Enz and G. C. Temes, "Circuit techniques for reducing the effects of op-amp imperfections: Autozeroing, correlated double sampling, and chopper stabilization," *Proc. IEEE*, vol. 84, no. 11, 1996.

[9] 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.

[10] 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.

[11] 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.

[12] T. Keller, A. di Girolamo, Y. Chen, H. Ghiasi, D. Seo, Y. Lee, and T. Jang, "A 1 nJ·mmHg^2 RDC using a double-settling RC bridge and a common-mode cancelling series-parallel amplifier for implantable pressure and strain monitoring," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC)*, 2026.