Tracking charge injection for the 2 Vpp noise-shaping SAR ADC of Fu *et al.* (CICC 2026)

Full report: model, implementation and evidence for five proposed changes

Nithin Purushothama
nithinpuru.com/ns_sar

10 November 2026. Simulation only: nothing in this report was built in silicon.

Abstract

The noise-shaping SAR ADC of [1] reaches a 2 Vpp input range and 5.62 µVrms noise at 2.26 µW. It subtracts a 4-bit prediction on the bottom plates of a passive capacitor stack and resolves the amplified residue with a charge-injection DAC. As drawn in its timing diagram, that DAC runs its full search, 511 charge packets, in every conversion, although after the prediction a slow neural signal leaves a residue of only a few packets. This report builds a sample-by-sample model of the converter, calibrated to the paper's noise, power split and search, and uses it to test five changes. Moving only the charge each residue needs (P1, tracking injection) and rotating the DWA only when the coarse code changes (P5) lower the modelled power from 2.26 to 1.29 µW on a 2 Vpp sine and from 2.12 to 0.91 µW on a quiet neural input (central scenario; 0.66 to 1.56 µW across scenarios). The noise falls from 5.55 to 4.81 µVrms and the NEF from 12.7 to 7.2, and the dynamic-range figure of merit rises by 4.9 dB (2.6 to 6.3 dB), with no loss of SNDR. A stimulation flag (P2) keeps fast artifacts up to 1.8 Vpp inside the search range. A shared injection engine (P3) and a finer prediction with more passive gain (P4) were studied and do not pay off on their own. P1 is implemented as a Verilog controller that matches a bit-exact model on 18 000 vectors, and the packet noise factor comes from a SKY130 transistor-level run of the injection cell.

1 The converter of [1]

Fig. 1 shows the architecture (after Figs. 1(c) and 2 of [1]). The differential input is sampled on 16 unit capacitors per side ($C_{\text{in}} = 8.4\,\text{pF}$) whose bottom plates are driven by a tri-level DWA with the predicted coarse code D_{pred} (4 bits, from earlier outputs; predict-and-skip [4]). Stacking eight stages in series gives the residue a passive gain of 8,

$$V_{\text{stack}} = 8\,(V_{\text{in}} - V_{\text{pred}}). \tag{1}$$

A charge-injection DAC, the cell of [2], then resolves 9 bits directly on the stacked node: each cycle moves one packet $Q_p = C_{CS}V_{GS}$, and the search fires bursts of $2^8, 2^7, \ldots, 2^0$ packets with the sign set by the comparator (Fig. 3 of [1]),

$$V_{\text{SAR}} = \frac{Q_p}{C_{\text{node}}} \sum_{k=0}^{8} s_k 2^k, \qquad s_k \in \{-1, +1\}, \tag{2}$$

so a conversion always costs $\sum_k 2^k = 511$ packets. A 2nd-order fully dynamic capacitor-stacking-and-buffering integrator shapes the error; a comparator and SAR logic close the loop and the digital back end aligns D_{pred} and D_{SAR} into a 12-bit output. Measured [1]: 128 kS/s, 500 Hz band (OSR 128), 2.26 µW at 1.2 V (digital 38 %, PAS-assisted SPA 31 %, charge injectors 21 %, comparator 7 %, integrator 3 %), 5.62 µVrms input-referred noise, 95.2 dB SNDR at 2 Vpp, 101.6 dB DR, 0.08 mm^2 in 65 nm.

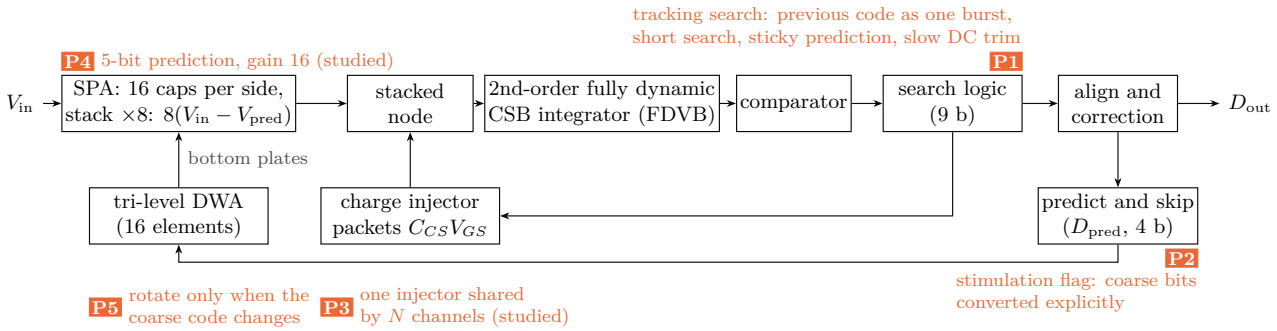


Figure 1: The converter of [1] (after its Figs. 1(c) and 2) with the proposed changes P1 to P5.

2 The model

Signal path. All quantities are referred to the input. With $x[n]$ the sampled input including its noise, $v_p[n]$ the analog prediction from the DWA and $e[n]$ the physical residue left on the node after the search, the search sees

$$y[n] = x[n] - v_p[n] + 2e[n-1] - e[n-2], \tag{3}$$

an ideal 2nd-order error-feedback loop (NTF $= (1 - z^{-1})^2$) that stands in for the CSB integrators. The digital output is $D_{\text{pred}}\,\Delta + D_{\text{SAR}}\,p$ with the coarse step $\Delta = 125\,\text{mV}$ and the packet $p = \Delta/511 = 245\,\mu\text{V}$ (the output LSB is two packets, because the sign-based sums are odd). The prediction is the previous output rounded to the coarse grid (an assumption; [1] defers the rule to [4]).

Noise. Three white sources: input sampling $\sqrt{2kT/C_{\text{in}}} = 31.4\,\mu\text{V}$ per sample; packet noise $\sqrt{\alpha kTC_{CS}}$ per packet referred through the node capacitance and the gain ($C_{CS} = 1.47\,\text{fF}$ for a $0.3\,\text{pF}$ node, $\alpha = 1.89$ from Sec. 3.1); and the rest, which the model cannot attribute, sized so that the paper's design hits $5.62\,\mu\text{Vrms}$ ($45.1\,\mu\text{V}$ per sample) and kept fixed when a proposal changes the design. Comparator noise ($0.4\,\text{mVrms}$ at the node) acts on decisions and is shaped like the quantization error. Because the loop measures the residue on the node, which contains the charge the DAC really injected, while the digital code assumes ideal packets, the packet noise reaches the output unshaped: $D_{\text{out}} = x - \text{NTF} \cdot q - n_{\text{DAC}}$.

Energy. Calibrated at a $2\,\text{Vpp}$ sine to $2.26\,\mu\text{W}$ with the Fig. 5 split. The paper does not say how each block scales with activity, so three scenarios bound it: injector energy proportional to packets with none (central, optimistic) or half (pessimistic) of it static; 0 (pessimistic), 50 % (central) or 80 % (optimistic) of the digital energy proportional to injector cycles; half of the SPA energy proportional to DWA toggles; comparator energy proportional to decisions.

Calibration and checks. The paper's design reproduces 511 packets and 9 decisions per sample, $5.55\,\mu\text{Vrms}$ simulated ($5.62\,\mu\text{Vrms}$ expected) and $2.26\,\mu\text{W}$ in all scenarios. Seven unit tests cover the calibration, the search counts, the unshaped DAC noise, P2 and P5. The model's SNDR (101.8 dB) is set by its noise floor only. The measured SNDR is lower: in Fig. 6(c) of [1] the points follow the input level plus about 99 dB at every level, while $5.62\,\mu\text{Vrms}$ predicts plus 105 dB (Fig. 2); I therefore compare SNDR only within the model and use the shorted-input noise as the reference.

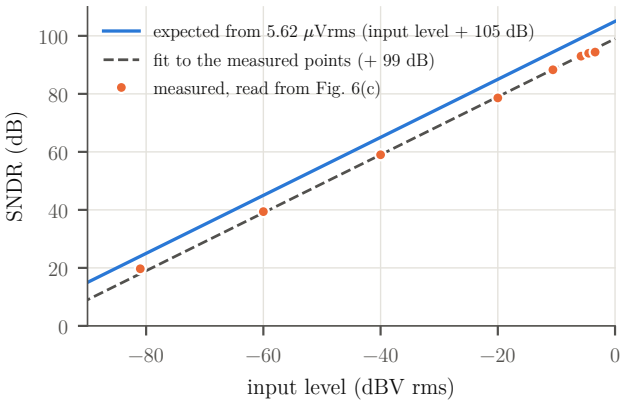


Figure 2: SNDR against input level: points read from Fig. 6(c) of [1], the line expected from 5.62 µVrms, and a fit.

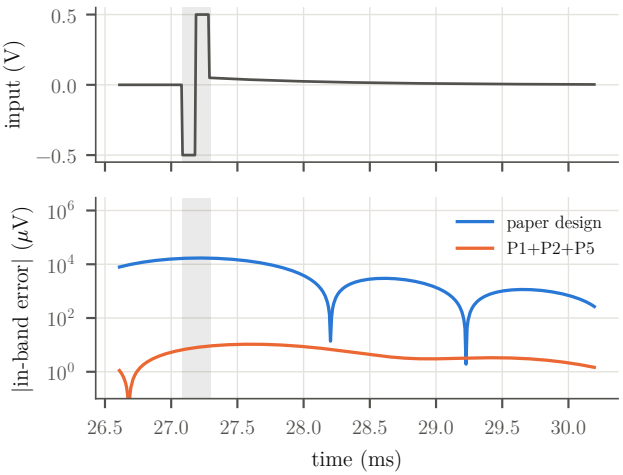


Figure 3: A 1 Vpp biphasic artifact with one-sample edges (top) and the in-band error of the paper design and of P1+P2+P5 (bottom, zero-phase 500 Hz low-pass). Grey: the flagged pulse window.

3 Where the energy and the noise go

The search costs 511 packets per sample whatever the residue, about 65 million packets per second; the paper's power split implies about 7.26 fJ per packet for the injectors alone. With predict-and-skip the residue of a 2 Vpp, 216 Hz sine needs on average far fewer packets (Fig. 4), and a quiet neural input needs almost none. On the noise side, the 8.4 pF input explains 2.78 µVrms (24 % of the noise power) if 8.4 pF is per side, 3.93 µVrms (49 %) if it is the total. The packets explain part of the rest (Fig. 5): with $\alpha = 1.89$ they contribute 3.99 to 1.55 µVrms (50 to 8 % of the noise power) for a node of 0.15 to 1 pF.

3.1 SKY130 check of the injection cell

A transistor-level SKY130 version of the cell (tt, 27 °C, 1.8 V; follower 2/0.5 µm, 50 MHz phases) was run for $C_{CS} = 2, 5, 10$ fF, $I_b = 0.5, 1, 2$ µA and node capacitances of 0.2 and 0.5 pF. The phase-1 noise was integrated from a small-signal noise analysis (self-test against kT/C within 0.3 %) and the phase-2 charge sharing added. The packet charge noise is $\sqrt{\alpha kTC_{CS}}$ with $\alpha = 1.62$ to 2.12 (median 1.89, used as the central value), the energy per packet is 12.4 to 71.4 fJ, dominated by the follower bias (so the bias should be gated between bursts), and the packet charge changes by 1.46 to 8.79 % over a 0.6 V node swing with the bulk at ground (2.2 % with a source-tied bulk; a smooth, mostly linear dependence). These are SKY130 estimates, not the paper's 65 nm cell.

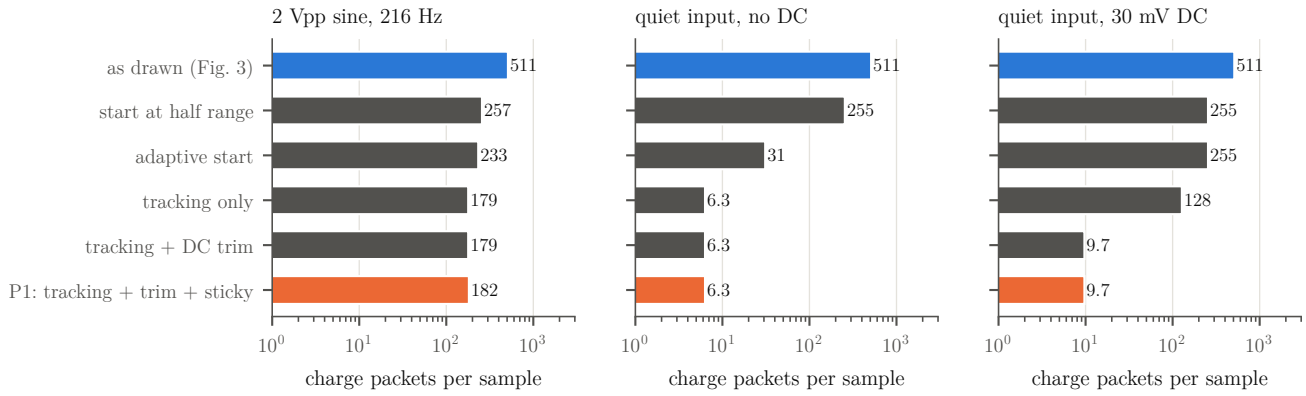


Figure 4: Charge packets per sample for six search orders and three inputs (model, 2^{16} samples). Blue: as drawn in Fig. 3 of [1]; orange: P1.

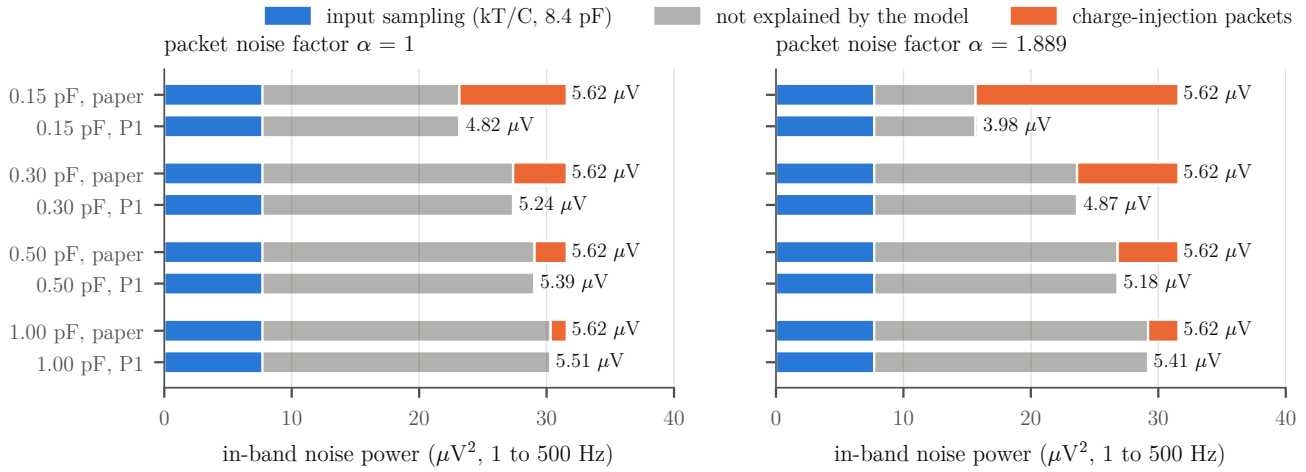


Figure 5: In-band noise power of the paper design and of P1 on a quiet input, split into input sampling, the part the model cannot attribute, and charge-injection packets, for four node capacitances and two packet noise factors.

4 Proposals and results

4.1 P1: tracking charge injection

1. Inject the previous code as one burst, shifted by one full range (± 511 packets) when D_{pred} moved and by the DC-trim change: no comparison is needed for this burst.
2. Start a sign-based search at the smallest burst 2^k whose bracket $2^{k+1}-1$ covers the previous remainder plus a margin of 2 packets.
3. If every decision had the same sign, the residue may lie outside the bracket: inject growing bursts $2, 4, 8, \ldots$ in that direction until the comparator flips, then refine below the last burst.
4. Keep the coarse code until the output moves 0.58 coarse steps away (sticky prediction), and move a slow DC-trim capacitor (step $\Delta/16$) towards the mean fine code while the input is quiet.

The final residue stays within one packet, so the loop is unchanged. Packets fall from 511 to 182 on the sine, 6.3 on a quiet input and 11.6 with 123 mV of electrode DC (11.0 on average over random DC within ± 0.3 V); decisions fall from 9 to 5.5 and 2.5. The controller is written in Verilog-2001 (one packet per clock) and checked with Icarus Verilog against a bit-exact Python mirror on 18 000 vectors from the model's residue sequences, with 0 mismatches and every final residue within one packet (Table 1).

Table 1: P1 in Verilog: mean clock cycles, packets and decisions per conversion (18 000 vectors).

input	search	cycles	packets	decisions
2 Vpp sine	as drawn / tracking	521 / 186	511.0 / 179.5	9.0 / 5.6
quiet input	as drawn / tracking	521 / 9.6	511.0 / 5.8	9.0 / 2.8
quiet, 123 mV DC (no trim)	as drawn / tracking	521 / 16.4	511.0 / 12.3	9.0 / 3.0

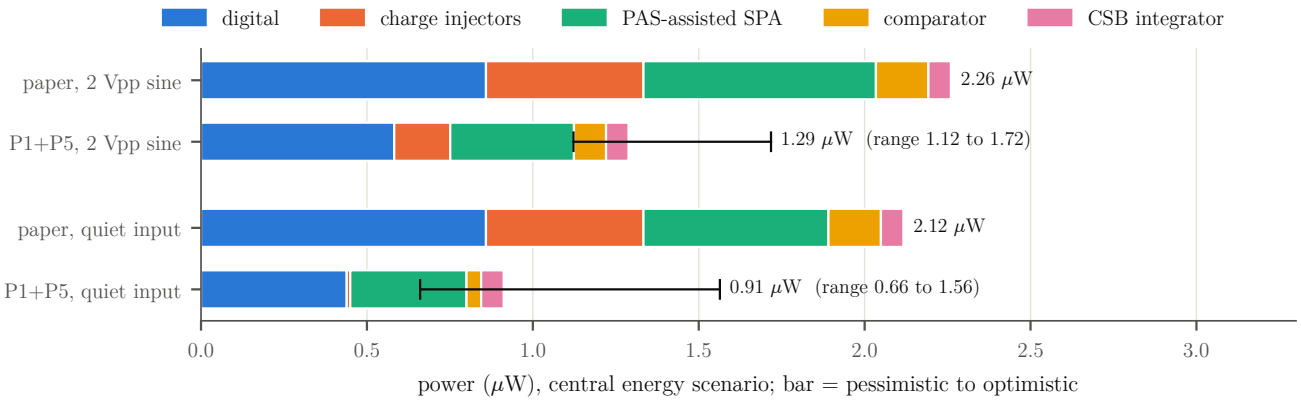


Figure 6: Modelled power by block, paper design against P1+P5 (central scenario; the line spans the pessimistic to optimistic totals).

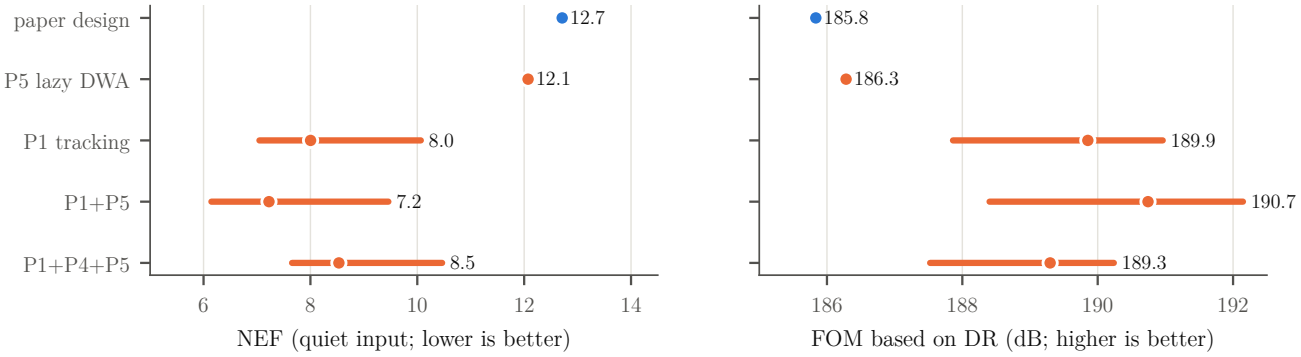


Figure 7: NEF and the dynamic-range figure of merit on a quiet input for each design: dot central scenario, bar pessimistic to optimistic.

4.2 P5: lazy DWA

While the coarse code holds, the same unit capacitors stay selected; the DWA pointer advances only when the code changes. Toggles fall from 6.76 to 0.43 per sample on the sine and to zero on a quiet input. Element mismatch is still averaged over code changes: the model's SNDR is 101.9, 101.8 and 101.6 dB for 0, 0.1 and 0.3 % unit mismatch with the paper's rotation, and 102.6, 102.6 and 102.3 dB with P1 and P5.

4.3 Combined result

Figs. 6 and 7 and Table 2 give the combined effect. Most of the saving comes from the injectors and the injection clock (P1); P5 removes most of the toggle-dependent SPA energy.

4.4 P2: stimulation-proof prediction

The closed-loop system knows when it stimulates. Inside a flagged window (pulse plus two samples) the coarse bits are converted explicitly on the bottom plates (four decisions) instead of being predicted, and the fine search starts at half range. Biphasic pulses at 130 Hz with 100 µs phases, one-sample edges and a 1 ms polarisation tail were added to a quiet input (33 pulses per run). Table 3 lists out-of-range samples and the in-band error; Fig. 3 shows one pulse. The paper's own in-vitro artifacts were 150 mVpp (its Fig. 7); their edges are not given.

4.5 P3: one injection engine shared by several channels

With P1 a quiet channel needs few packets, so one fast injector, comparator and search logic could serve several channels while each keeps its stack and integrators. Between pulses eight channels need 180 packets per period on average, but artifacts reach all channels together: their decaying tails move every coarse code at about the same time, so the 99.9th percentile is 1993 and the worst pulse 2223 (Fig. 8). The engine would need about four times today's packet rate, or a local fallback during pulses. Verdict: **partial**.

Table 2: Paper design against the proposals (model; central scenario, pessimistic to optimistic in brackets; quiet input with 123 mV electrode DC).

	paper design	P1	P5	P1+P5
packets per sample, sine / quiet	511 / 511	182 / 11.6	511 / 511	182 / 11.6
power, sine (µW)	2.26	1.62 (1.45 to 2.05)	1.93	1.29 (1.12 to 1.72)
power, quiet (µW)	2.12	1.12 (0.87 to 1.77)	1.91	0.91 (0.66 to 1.56)
noise (µVrms)	5.55	4.81	5.55	4.81
NEF, quiet	12.7	8.0	12.1	7.2 (6.1 to 9.5)
FOM from DR, quiet (dB)	185.8	189.9	186.3	190.7
SNDR, model (dB)	101.8	102.7	101.8	102.7

Table 3: Out-of-range samples per pulse and in-band error (µVrms) with fast artifacts (model).

artifact (Vpp)	paper design	P1+P5	P1+P2+P5
0.15	3 / 54.9	1 / 4.4	0 / 4.4
0.50	14 / 1 492	5.5 / 675	0 / 4.4
1.00	22 / 5 826	12 / 2 984	0 / 4.5
1.80	28 / 26 323	17 / 9 174	0 / 4.5

4.6 P4: 5-bit prediction with a passive gain of 16

If the 16 tri-level elements per side form 33 levels, the residue halves and 16 stacking stages fit the swing, so everything after the stack is divided by 16. The noise falls from 5.55 to 4.99 µVrms (4.80 with P1 and P5), but assuming the fixed SPA energy doubles with the stages, the power rises to 2.62 µW. Its prediction DAC sits before the gain, so its errors must hold to the full resolution (DWA shapes them: 101.8, 101.6 and 101.3 dB at 0, 0.1 and 0.3 %). Verdict: **no win** alone; P1 obtains most of the noise benefit without it.

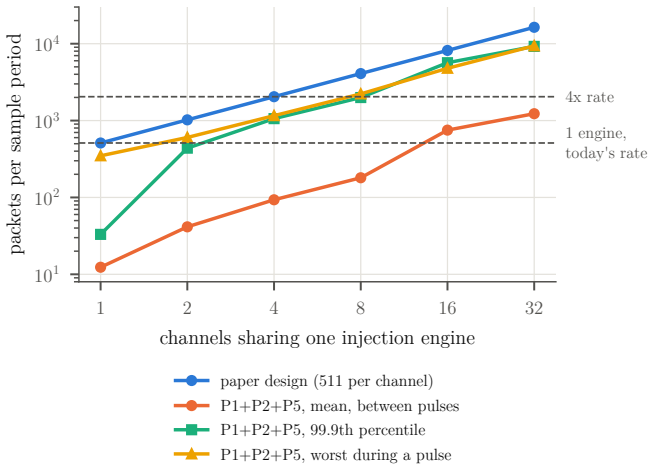


Figure 8: Packets per sample period that one shared injection engine must inject for N channels with common stimulation (model, 32 channels with random electrode DC and artifact gain).

Table 4: Verdicts.

	what it does	verdict	closest prior work
P1	previous code as one burst, short search, growth fallback, sticky prediction, DC trim	**win**	LSB-first SA [5], bypass window [6]; array-assisted charge-injection SAR [3]
P5	rotate the DWA only when the coarse code changes	**win**	prediction-assisted mismatch-shaping DEM [7] (not checked in detail)
P2	explicit coarse bits inside known stimulation windows	**win** (robustness)	artifact cancellation before amplification [8,9]; predictive autoranging [10]
P3	one injector, comparator and search logic for N channels	**partial**	time-multiplexed front ends share whole ADCs
P4	5-bit prediction, passive gain 16	**no win** alone	delta amplification [11]

5 Sensitivity and limits

Node capacitance. The noise gain of P1 is 3.0 dB for a 0.15 pF node, 1.2 dB for 0.3 pF and about 0.3 dB for 1 pF; the power gain does not depend on it. **Energy scaling.** If the injector bias stays on and the digital power does not follow the injection clock (pessimistic scenario), P1+P5 still saves 26 % of the quiet-input power (69 % optimistic). **Coarse/fine match.** One coarse step must equal 511 packets. With a symmetric coarse range the model loses little at 0.1 % mismatch (101.8 dB) and 1.4 dB at 0.3 % (100.6 dB); with an asymmetric range (−8 to +7) the same mismatch costs far more, so the range should stay symmetric or the ratio be calibrated. **Model limits.** An ideal loop replaces the CSB integrators and their buffer; reference and supply noise, settling, charge injection of the switches and the stacking parasitics are not modelled; the prediction rule is assumed; the energy model is a calibrated split, not a circuit simulation.

6 Next steps

A SKY130 transistor-level version of the stack, the injection cell and the P1 controller would replace the energy scenarios with simulated energy; on the authors' chip, the two numbers that decide the noise gain (the node capacitance and the cell noise) could be read directly, and the injection-clock share of the digital power measured by running the search at a fixed shorter schedule.

A The fully dynamic buffer as a dynamic follower

The FDVB of [1] replaces the follower's bias current by a pre-charged capacitor, the integration capacitor itself. Without a bias, the current falls roughly as $1/t$ once the devices enter subthreshold, and the output keeps creeping logarithmically: there is no fixed time constant (Fig. 9; EKV-style model, not fitted to a PDK). The creep is the same for every input, so the paper's odd/even chopping removes it; the gain stays near one only if the bulks follow the sources (about $1/n \approx 0.77$ otherwise). This is why a sizing method for this buffer has to integrate the device current rather than assume a constant g_m.

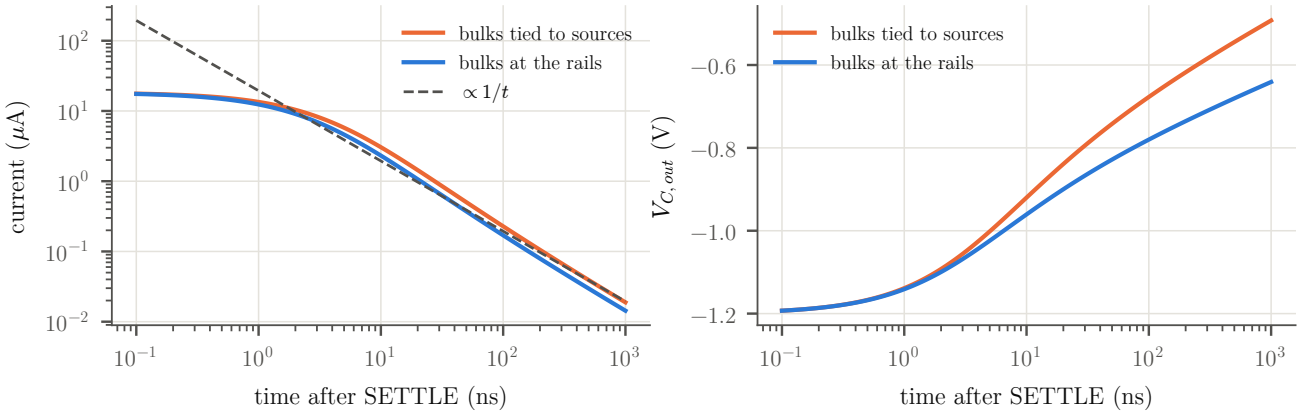


Figure 9: Dynamic complementary follower settling onto its pre-charged capacitor: current (left, with a $1/t$ guide) and the capacitor voltage (right), with bulks tied to sources and at the rails.

B Reproduce

In the author's tracking repository, folder `universities/EPFL/Choo/`:

- `model/`: `python -m pytest -q tests`; `python results.py`; `python figs.py`;
- `model/rtl/`: `python rtl/run_rtl.py` (needs Icarus Verilog);
- `model/spice/`: `python cidac_cell.py` (needs ngspice and the SKY130 models);
- `report/`: `python make_tex.py`, then `latexmk -pdf summary.tex main.tex`.

Every number in this report is a macro written by `make_tex.py` from the results files.

References

[1] Y. Fu *et al.*, "A 0.08mm^2, 2Vpp-input-range, 95dB-SNDR, 178.5dB-FOMs noise-shaping SAR ADC for closed-loop neural interfaces," in *Proc. IEEE CICC*, 2026, doi:10.1109/CICC65509.2026.11509539.

[2] K. D. Choo, J. Bell, M. P. Flynn, "Area-efficient 1GS/s 6b SAR ADC with charge-injection-cell-based DAC," in *IEEE ISSCC*, 2016.

[3] K. D. Choo *et al.*, "Energy-efficient low-noise CMOS image sensor with capacitor array-assisted charge-injection SAR ADC for motion-triggered low-power IoT applications," in *IEEE ISSCC*, 2019.

[4] H. Zhao *et al.*, "A fully dynamic noise-shaping SAR ADC achieving 120dB SNDR and 189dB FoMs in 1kHz BW," in *IEEE ISSCC*, 2025.

[5] F. M. Yaul, A. P. Chandrakasan, "A 10 bit SAR ADC with data-dependent energy reduction using LSB-first successive approximation," *IEEE JSSC*, vol. 49, no. 12, 2014.

[6] G.-Y. Huang *et al.*, "A 1-µW 10-bit 200-kS/s SAR ADC with a bypass window for biomedical applications," *IEEE JSSC*, 2012.

[7] Y. Han *et al.*, "An 81.7MΩ-input-impedance 179.5dB-FOMSNDR 1.8Vpp-input-range noise-shaping-SAR-based sensing frontend with dynamic input-impedance boosting and prediction-assisted mismatch-shaping-DEM," in *Proc. IEEE CICC*, 2025.

[8] A. E. Mendrela *et al.*, "A bidirectional neural interface circuit with active stimulation artifact cancellation and cross-channel common-mode noise suppression," *IEEE JSSC*, vol. 51, no. 4, 2016.

[9] J. P. Uehlin *et al.*, "A single-chip bidirectional neural interface with high-voltage stimulation and adaptive artifact cancellation in standard CMOS," *IEEE JSSC*, vol. 55, no. 7, 2020.

[10] C. Kim *et al.*, "Sub-µVrms-noise sub-µW/channel ADC-direct neural recording with 200-mV/ms transient recovery through predictive digital autoranging," *IEEE JSSC*, vol. 53, no. 11, 2018.

[11] K. Jeong *et al.*, "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 ISSCC*, 2025.