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

Three changes that make each conversion do only the work its residue needs

Nithin Purushothama nithinpurushothama@gmail.com interactive version: nithinpuru.com/ns_sar 10 November 2026

Abstract. In the converter of [1], predict-and-skip leaves only a small residue for the passive gain, but the charge-injection DAC still runs its full 9-bit search in every conversion (Fig. 3 of [1]): 511 charge packets per sample, even when the residue is only a few packets. The charge injectors and the digital logic take 59 % of the 2.26 µW, and each packet adds thermal noise that the noise-shaping loop does not suppress. I propose three changes so that the work per sample follows the residue. P1 starts the search from the previous code, P5 rotates the DWA only when the coarse code changes, and P2 converts the coarse bits directly during a known stimulation pulse. I tested them in a sample-by-sample model calibrated to the published noise and power split. P1 cuts the packets per sample from 511 to 182 on a 2 Vpp sine and to 11.6 on a quiet neural input. With P5 added, the power drops from 2.26 to 1.29 µW on the sine and from 2.12 to 0.91 µW on the quiet input, the noise from 5.55 to 4.81 µVrms and the NEF from 12.7 to 7.2, with no loss of SNDR. The dynamic-range figure of merit improves by 4.9 dB (2.6 to 6.3 dB, depending on how the digital power scales). P2 keeps fast artifacts up to 1.8 Vpp inside the search range. The tracking controller is written in Verilog and matches the model bit for bit, and the packet noise is taken from a SKY130 transistor-level simulation of the injection cell. All results are from simulation.

What is new

P1	**Tracking charge injection.** Inject the previous code in one burst, then search only the remainder.	Packets per sample 511 to 182 (2 Vpp sine) and 11.6 (quiet input); 521 to 186 clock cycles per conversion; noise −1.2 dB.
P5	**Lazy DWA.** Rotate the unit capacitors only when the coarse code changes.	DWA switching 6.76 to 0.43 toggles per sample on the sine, none on a quiet input; no SNDR loss at 0.3 % unit mismatch.
P2	**Stimulation-proof prediction.** In a known pulse window, convert the coarse bits instead of predicting them.	Out-of-range samples per pulse 3 to 28 without it, 0 with it, for 0.15 to 1.8 Vpp artifacts with fast edges.

Together: 2.26 to 1.29 µW on a 2 Vpp sine, 2.12 to 0.91 µW on a quiet input, NEF 12.7 to 7.2, FOM +4.9 dB, SNDR unchanged (model, central case).

1 Where the work goes

The input is sampled on 16 capacitors per side whose bottom plates carry a 4-bit prediction; stacking them gives $8(V_{\text{in}} - V_{\text{pred}})$ without an amplifier, and the charge-injection DAC [2] resolves 9 bits on the stacked node with packets of charge $C_{CS}V_{GS}$ (Fig. 1). As drawn in Fig. 3 of [1], the search always fires bursts of 256, 128, …, 1 packets, so every sample costs 511 packets, about 65 million per second, and most of the charge moved by the large bursts is taken back by the next ones. After the prediction, a slow neural signal leaves a residue of only a few packets. The injectors and the digital logic use 21 % and 38 % of the power, against 3 % for the integrator. On the noise side, the 8.4 pF input explains only 2.78 of the 5.62 µVrms (24 % of the noise power), and every packet adds kT/C-type noise that reaches the output unshaped: the loop sees the charge the DAC really injected, while the output code assumes ideal packets.

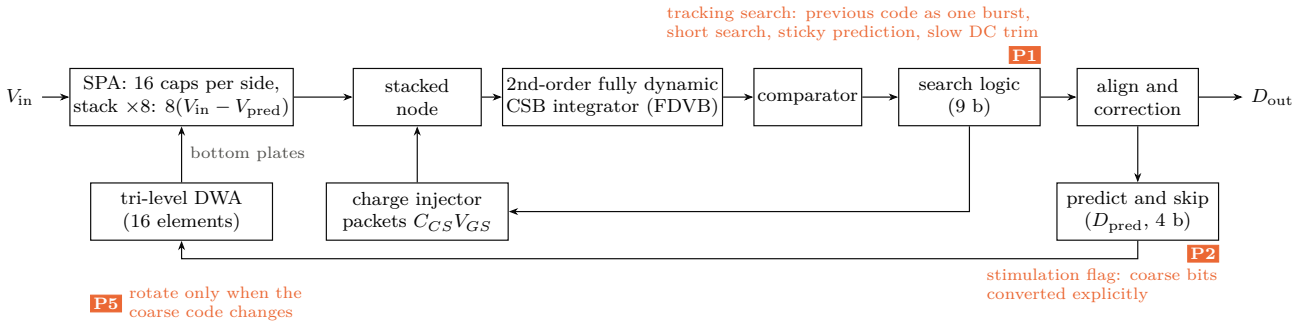


Figure 1: The converter of [1] (after its Figs. 1(c) and 2) with the three proposed changes.

2 How each change works

P1, tracking charge injection. The search first injects the previous sample's code in one burst; this needs no comparisons because its sign is known (shifted by one full range if the coarse code has moved). A short sign-based search then resolves what is left, starting from the smallest burst that covers the previous remainder plus two packets. If all its decisions agree, the residue lies outside that bracket, and the logic grows the bursts until the

Table 1: The published chip, the model of it, and the model with P1, P5 and P2. Power ranges: how much of the digital power follows the injection clock is not published, so it is varied.

	measured [1]	paper design, model	with P1, P5, P2, model
charge packets per sample: 2 Vpp sine / quiet input		511 / 511	182 / 11.6
power, 2 Vpp sine (µW)	2.26	2.26	1.29 (1.12 to 1.72)
power, quiet input with electrode offset (µW)		2.12	0.91 (0.66 to 1.56)
input-referred noise, 1 to 500 Hz (µVrms)	5.62	5.55	4.81
NEF, quiet input[a]	14.61	12.7	7.2 (6.1 to 9.5)
dynamic-range figure of merit, quiet input (dB)		185.8	190.7
SNDR at 2 Vpp (dB)[b]	95.2	101.8	102.7
out-of-range samples per pulse, 1.8 Vpp fast artifact		28	0

[a] Measured: Table I of [1]. Model: $I_{\text{tot}} = P/1.2\,\text{V}$, which gives 13.3 for the measured 5.62 µVrms and 2.26 µW. [b] The model's SNDR is set by its noise floor only; the measured SNDR is lower at every input level (Fig. 6(c) of [1]), so SNDR is compared within the model.

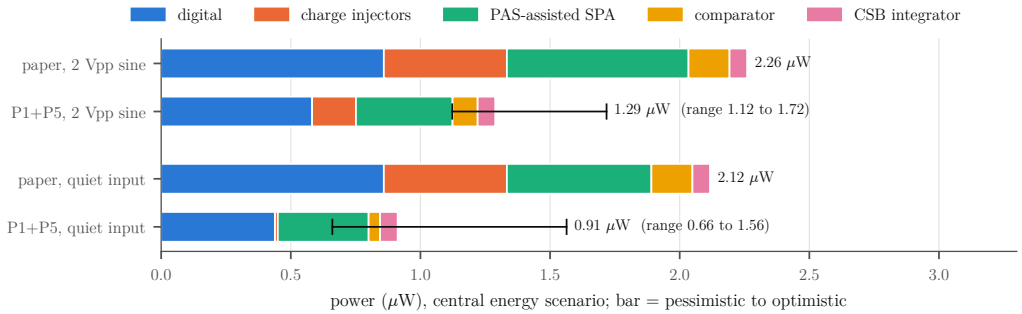


Figure 2: Modelled power by block, paper design against P1 with P5, for a 2 Vpp sine and for a quiet input with electrode offset. The line spans the range of the digital-power assumption.

comparator flips. The final residue is still within one packet, so the noise-shaping loop is unchanged. Two small additions keep quiet inputs cheap: the coarse code changes only when the output moves 0.58 of a step away, so it does not toggle around a threshold, and a slow trim capacitor (1/16 of a step) absorbs the fine part of the electrode offset. In Verilog the controller needs 186 clock cycles per conversion on the sine instead of 521, and about 9.6 on a quiet input; 18 000 test vectors match the model bit for bit. Data-dependent searches already save switching in capacitor DACs [4,5], and Choo *et al.* cut injection steps with a capacitor array [3]. Here the saving follows the residue, and in a charge-injection DAC both energy and noise scale with the charge moved.

P5, lazy DWA. The tri-level DWA rotates its 16 unit capacitors every sample, even when the coarse code does not change. Rotating only when the code changes removes almost all of this switching (6.76 to 0.43 toggles per sample on the sine, none on a quiet input). In the model this costs no SNDR at 0.1 and 0.3 % unit mismatch, since the mismatch is still averaged over code changes. A related idea is the prediction-assisted mismatch shaping of [6].

P2, stimulation-proof prediction. A closed-loop system knows when it stimulates. In a flagged pulse window the coarse bits are found by a four-step search on the bottom plates, not predicted, so a fast edge cannot push the residue out of range. With one-sample edges, the modelled paper design leaves the range 3 to 28 times per pulse for 0.15 to 1.8 Vpp artifacts, with an in-band error of 26 mVrms at 1.8 Vpp; with P2 it never does, and the error stays at the noise floor (4.5 µVrms). The in-vitro test of [1] used 150 mVpp artifacts of unstated edge speed. Earlier work cancels artifacts before amplification [7,8] or autoranges the ADC [9]; P2 reuses the prediction path.

3 How this was checked, and what is assumed

The model follows the converter sample by sample and reproduces the search of Fig. 3 of [1] (511 packets, 9 decisions), its 5.62 µVrms noise and its 2.26 µW power split. The packet noise comes from a SKY130 transistor-level run of the injection cell (1.62 to 2.12 times kT/C over 18 sizings). Two numbers not given in [1] are varied, not guessed: the capacitance at the injection node (0.3 pF in the central case; the noise gain is 3.0 dB at 0.15 pF and 0.3 dB at 1 pF) and how much of the digital power follows the injection clock (the ranges in Table 1). The prediction rule and an ideal second-order loop in place of the CSB integrators are assumptions. Two further ideas, P3 (one injector shared by many channels) and P4 (a finer prediction with more passive gain), do not pay off on their own; they are in the full report at nithinpuru.com/ns_sar, with the code and an interactive model.

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," CICC 2026. [2] K. D. Choo, J. Bell, M. P. Flynn, "Area-efficient 1GS/s 6b SAR ADC with charge-injection-cell-based DAC," ISSCC 2016. [3] K. D. Choo *et al.*, "Energy-efficient low-noise CMOS image sensor with capacitor array-assisted charge-injection SAR ADC," ISSCC 2019. [4] F. M. Yaul, A. P. Chandrakasan, "A 10 bit SAR ADC with data-dependent energy reduction using LSB-first successive approximation," JSSC 2014. [5] G.-Y. Huang *et al.*, "A 1-µW 10-bit 200-kS/s SAR ADC with a bypass window for biomedical applications," JSSC 2012. [6] Y. Han *et al.*, "An 81.7MΩ-input-impedance 179.5dB-FOMSNDR 1.8Vpp-input-range noise-shaping-SAR-based sensing frontend," CICC 2025. [7] A. E. Mendrela *et al.*, "A bidirectional neural interface circuit with active stimulation artifact cancellation," JSSC 2016. [8] J. P. Uehlin *et al.*, "A single-chip bidirectional neural interface with high-voltage stimulation and adaptive artifact cancellation," JSSC 2020. [9] C. Kim *et al.*, "Sub-µVrms-noise sub-µW/channel ADC-direct neural recording with 200-mV/ms transient recovery through predictive digital autoranging," JSSC 2018.