Rebuilding and Extending a Light-Adaptive Optoelectronic Neural Recording Mote: a Behavioural and SKY130 Study of the FB-MOTE

Nithin Purushothama
nithinpurushothama@gmail.com nithinpuru.com github.com/chennakeshavadasa

November 2026

Abstract

I rebuilt, block by block, the recording chain of the FB-MOTE, a 43 μm × 269 μm tetherless neural recording microsystem by Jiao *et al.* (Prof. Sunwoo Lee's group with Cornell co-authors) [1], as a Python behavioural model and, for the amplifier and the encoder oscillator, as SKY130 transistor-level replicas in ngspice. The model reproduces the measured bandwidth, noise density and oscillator period within 5.8 %, 8.3 % and 1.5 % RMS, and the optical-powering noise within 3.3 % using a single factor. I then tested nine ideas on the encoding, the optical pulse link, receiver-side decoding, noise and calibration. The strongest are a grid-tracking soft-decision decoder that needs 4-10 instead of 15 detected photons per pulse, a bandwidth clamp that removes about half of the noise aliasing at high light, and send-on-delta encoding that keeps spike detection with 68 % fewer samples. Several ideas are reported as negatives. All my results are simulation; the paper's numbers are measurements.

1 Introduction

Tetherless microscale recorders avoid the tissue damage of tethered electrodes. The FB-MOTE [1] forward-biases the CMOS bulk so that light both powers the chip and lowers transistor thresholds, lets an inverter-based amplifier absorb the extra light-generated current so that its bandwidth grows at constant integrated noise, encodes the signal as pseudo-chopped PWM, converts it to pulse-position modulation (PPM) and sends complementary pulse pairs through an AlGaAs micro-LED [2, 3, 4]. It works from 100 to 1200 μW/mm^2, with about 10 μV_{rms} integrated noise electrically and 14 to 15 μV_{rms} under light.

This report asks three questions: (i) can the chain be rebuilt from the paper alone, with each block checked against the measured curves; (ii) what limits the recording, the encoding and the optical link once the photovoltaic side is taken as given (I treat light only as an input that sets the bias current); and (iii) which changes, mostly on the encoding and receiver side, improve it. Section 2 describes the model, Section 3 validates it, Section 4 tests the ideas, Section 5 collects the negatives and Section 6 gives the outlook. Code, data and every number are in the accompanying repository (`demos/Molnar`).

2 The paper rebuilt

Neural signal. Synthetic LFP ($1/f$-shaped, 0.5-300 Hz, 150 μV_{rms}), three spike units (−60, −100, −150 μV peak at 8, 12, 20 Hz) and 6.3 μV_{rms} of biological noise, the difference between the in-vivo and open-air noise in Sec. IV-D of [1].

Amplifier. In subthreshold, g_m and the load conductance G_o are proportional to the amplifier current $I - I_o$ (I_o: current of the other, PTAT-biased blocks). Hence the corner $f_L = G_o/(2\pi C_s)$ and the noise density $\sqrt{S_v} \propto 1/\sqrt{g_m}$ give

$$f_L = k_L(I - I_o), \qquad \sqrt{S_v} = \frac{c}{\sqrt{I - I_o}}, \qquad v_{n,\mathrm{rms}} \approx \sqrt{S_v f_L} = c\sqrt{k_L}, \tag{1}$$

a constant. A joint fit to Fig. 11 gives $k_L = 11902\,\mathrm{Hz}/\mu\mathrm{A}$, $I_o = 0.119\,\mu$A, $c = 0.0908\,\mu\mathrm{V}\sqrt{\mu\mathrm{A}/\mathrm{Hz}}$, predicting 9.9 μV_{rms}. The time-domain amplifier adds white and $1/f$ noise (corner 50 Hz, assumed), a 2 Hz first-order high-pass and a first-order low-pass at $f_L(I)$.

Encoder and pulse generator. The oscillator period follows $T = AI^{-q}$ with $q = 0.335$ (+50 % current gives +14.5 % clock; the paper states about 15 %). Each period the amplifier output is sampled at mid-period;

the encoding pulse sits at $t_{\mathrm{ref}} + T(0.5 + o \pm gv)$ with the sign alternating (pseudo-chopping) and an encoder offset o. The time gain g follows Fig. 11 (top right). I assume a $\pm 1.5\,\mathrm{mV}$ full scale.

Optical channel. Each pulse yields Poisson(N_s) detected photons spread over 20 ns with 1 ns jitter; background (ambient, leaked 637 nm light, dark counts) arrives at λ_b counts/s; the detector bins at 10 ns. These are assumptions: the paper uses an APD and an oscilloscope and models the noise as additive white Gaussian (Sec. V).

Decoder (baseline). A re-implementation of Sec. V: per 100 ms batch, tanh-compressed histogram and Otsu threshold [7], non-maximum suppression ($\tau_s = 200\,\mathrm{ns}$, $\tau_p = 1\,\mu\mathrm{s}$), period from the shorth of lag-2 spacings, reference seeding from the folded onsets, and the brightest suprathreshold sample as the encoding pulse. Two documented deviations make it as favourable as possible: the reference grid re-anchors on each found pulse ($\pm 200\,\mathrm{ns}$), and the chopping parity is carried across batches by counting periods.

3 Baseline validation

Table 1 and Fig. 1 compare each block with the paper. Under light, the measured noise density is 1.43 times the electrical value at the same bandwidth at every level (Figs. 11-12 read together), so I model the optical excess as extra white noise proportional to the amplifier's own noise. Table 2 gives the integrated noise; the decoded chain is above Eq. (8) because a first-order roll-off integrates more noise than a brick wall, plus the assumed flicker noise.

Block check	RMS error vs paper
Amplifier bandwidth vs current (Fig. 11)	5.8 %
Input noise density vs current (Fig. 11)	8.3 %
Integrated noise, Eq. (8) (Fig. 11)	8.9 %
Oscillator period (Fig. 11)	1.5 %
Optical noise density (Fig. 12)	2.3 %
Optical integrated noise (Fig. 12)	3.3 %
Optical oscillator period (Fig. 12)	10.8 %

Table 1: Block-by-block agreement with the measured curves (digitised, about 3 % reading error).

I (μA)	paper IRN	model, Eq. (8)	model, decoded chain
0.2	8.6	9.9	10.5
0.4	9.1	9.9	10.8
0.6	9.1	9.9	11.8
0.8	9.2	9.9	12.0
1.0	9.4	9.9	11.9
1.2	9.7	9.9	11.7

Table 2: Integrated input-referred noise ($\mu\mathrm{V_{rms}}$), electrical bias. Optical: model 14.3-17.1 versus paper 13.5-14.7.

Fig. 2 repeats the paper's Fig. 13 experiment. Spike detection improves with light, the power-adaptive behaviour the paper reports: F1 0.52 / 0.71 / 0.79 / 0.81 at 100 / 480 / 900 / 1110 $\mu\mathrm{W/mm^2}$ (perfect link: 0.52 / 0.72 / 0.80 / 0.81).

Transistor level. A SKY130 half-circuit of the amplifier (LVT devices, NMOS bulk at VDD and PMOS bulk at ground, i.e. forward-biased bulk) was swept over the paper's VDD range, 0.317-0.43 V [14]. The current rises from 30 to 281 nA and the corner from 1884 to 12589 Hz (paper, optical: 1.9 to 12.2 kHz) while the integrated noise stays within 3.5 % (Fig. 3). The gain is 16.5-18.0 dB (paper: 25 dB). A transistor-level relaxation oscillator with a bias that tracks VDD gives 119.6 μs at 0.317 V (paper: 118 μs) but tracks the supply more strongly than the chip (period ratio 3.33 versus 1.83).

4 Ideas, baseline versus proposed

4.1 Idea 1: period-tracked receiver gain

Method. The decoder measures every period, so it can estimate light. I divide the pulse delay by the measured period and correct the residual gain change with a 5-point calibration $g(T)$ (1 % error per point). Light flickers

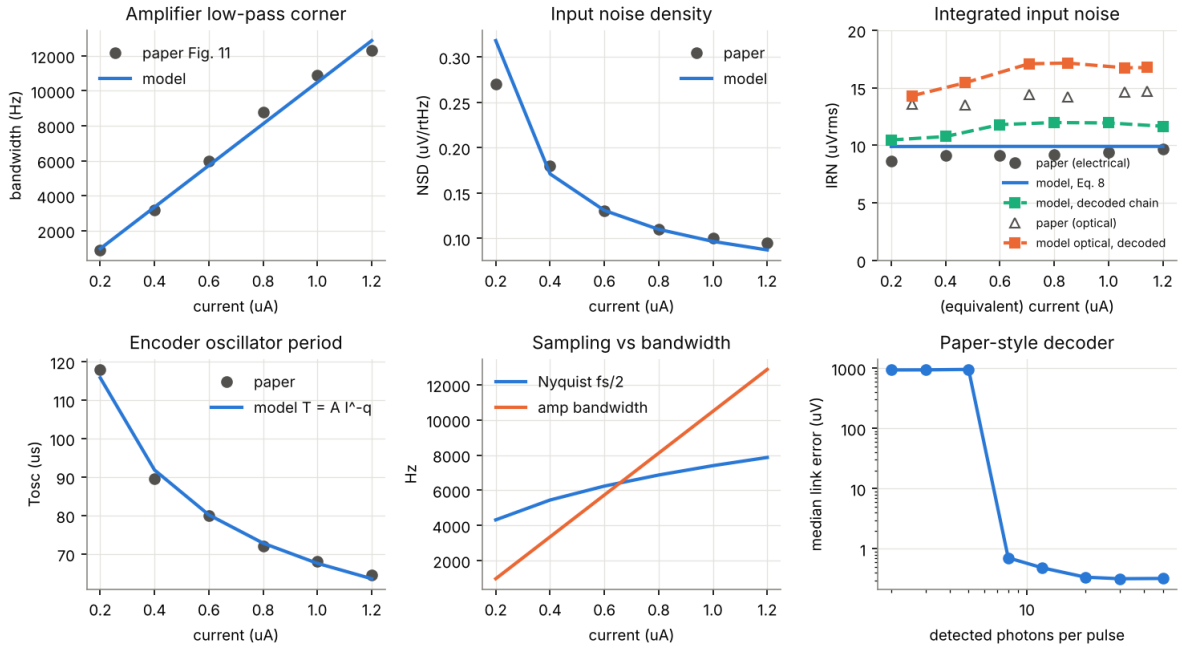

Figure 1: Paper (gray) versus model. Bottom middle: the encoder's Nyquist frequency falls below the amplifier bandwidth above about 0.65 μA. Bottom right: the paper-style decoder versus detected photons per pulse.

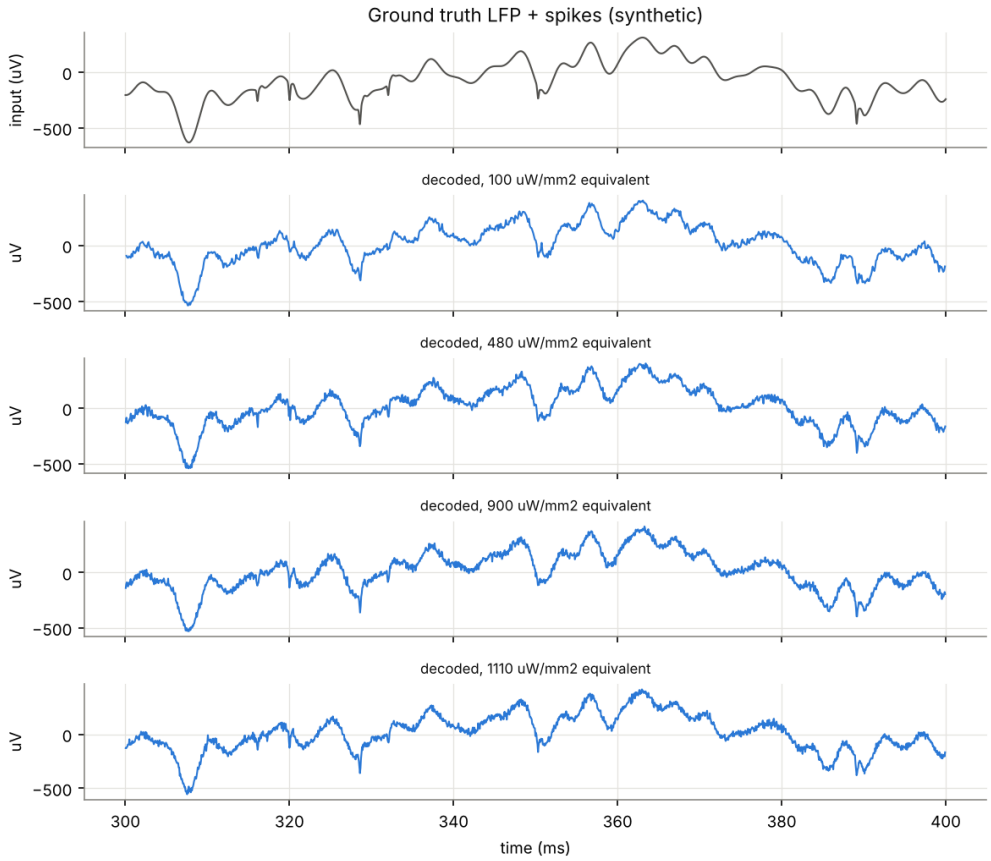

Figure 2: Analogue of the paper's Fig. 13: synthetic ground truth and decoded output at four irradiances (30 photons per pulse, paper-style decoder).

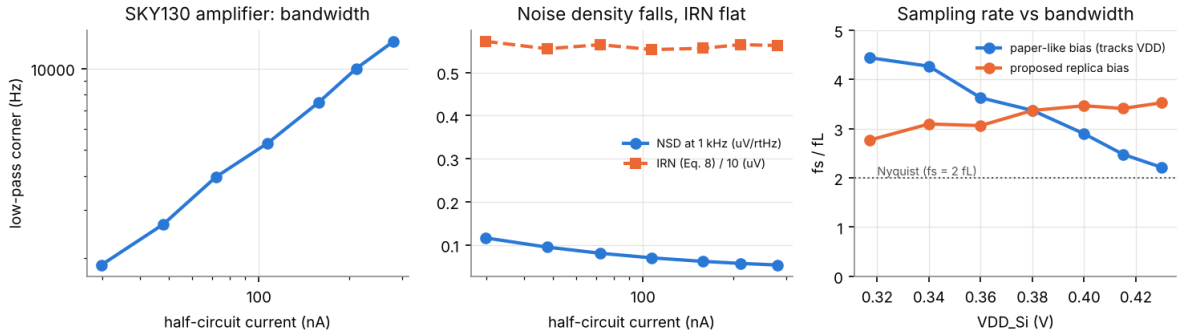

Figure 3: SKY130 replicas. Left, middle: amplifier corner rises with current, noise density falls, integrated noise stays flat. Right: sampling rate over amplifier corner for a supply-tracking bias and for the proposed replica bias (Idea 3).

between 420 and 980 μW/mm^2 at 1.5 Hz. *Result* (Table 3, Fig. 4): the gain error falls from 4.18 % to 1.14 % RMS. Spike F1 does not change: this idea protects amplitudes (LFP), not spike detection.

Receiver	gain err. RMS (%)	peak to peak (%)	spike F1
Fixed time scale (paper-style)	4.18	11.25	0.838
Ratiometric, fixed gain	2.18	6.56	0.830
Period-tracked gain (proposed)	1.14	2.62	0.834
Oracle gain	0.00	0.01	0.834

Table 3: Idea 1 under light flicker.

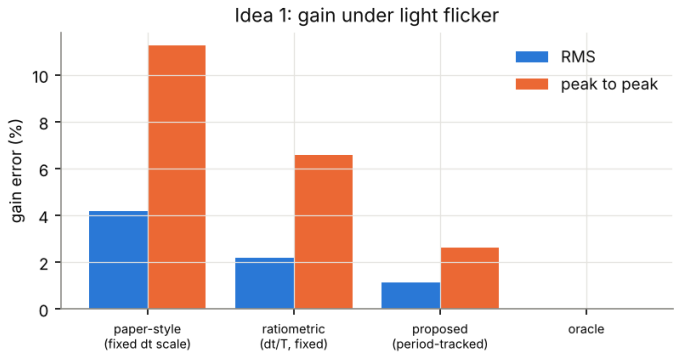

Figure 4: Idea 1: gain error under light flicker for four receivers.

4.2 Idea 2: grid-tracking soft-decision decoder

Method. A phase-locked loop tracks the reference pulses (acquired from exact triplets t, $t+T$, $t+2T$, which are robust to a changing period), so the encoding pulse is searched only inside its period. Every photon cluster is scored by $c\ln(1+N_s/\lambda_b w)$ plus the log of a Gaussian continuity prior on its position; if nothing beats the "pulse lost" hypothesis the sample is imputed from the prediction. This is the soft-decision idea of deep-space PPM receivers [5, 6] applied to analog-time PPM with a signal prior. *Result* (Fig. 5): to reach 95 % of the perfect-link F1 the paper-style decoder needs 15 photons per pulse at every background; the proposed one needs 4, 8 and 10 at 2×10^4, 2×10^5 and 10^6 counts/s. The paper-style decoder is non-monotonic (its Otsu threshold flips at some photon counts), and at 10^6 counts/s and 6 photons it scores higher than mine.

4.3 Idea 3: keep the sampling rate above twice the bandwidth

Method. At high light the chip samples near 14 kHz with about 12 kHz bandwidth, so amplifier noise between $f_s/2$ and f_L aliases into the spike band. Two fixes: clamp $f_L \le 0.4 f_s$ (switch extra C_s in), or bias the oscillator from a scaled replica of the amplifier's self-biased inverter so f_s follows f_L. The SKY130 replica bias keeps f_s/f_L within 2.76-3.52 (spread 1.27×) versus 4.44 down to 2.21 for the supply-tracking bias (Table 4). *Result*

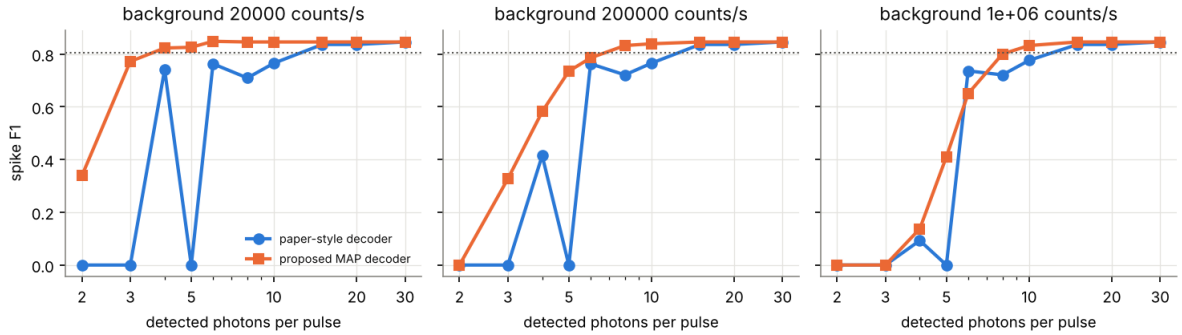

Figure 5: Idea 2: spike F1 versus detected photons per pulse at three background rates. Dotted: 95 % of the perfect-link F1.

(Table 5, Fig. 6): at 1.2 μA the spike-band noise falls from 9.43 to 6.97 μV with the clamp at no pulse cost, and to 6.48 μV ratio-locked at 2.65× the pulse rate. Ratio-locking alone fails at the lowest light (F1 0.44 to 0.11) and needs a minimum rate.

VDD (V)	supply-tracking T (μs)	f_s/f_L	replica T (μs)	f_s/f_L
0.317	119.6	4.44	192.1	2.76
0.34	88.1	4.27	121.5	3.09
0.36	69.3	3.63	82.1	3.06
0.38	55.9	3.37	55.9	3.37
0.4	46.1	2.89	38.5	3.46
0.415	40.4	2.48	29.3	3.41
0.43	35.9	2.21	22.5	3.52

Table 4: SKY130 oscillator, two bias schemes.

I (μA)	paper-like	clamp	ratio-locked	pulse-rate ratio
0.2	10.62 (0.77)	10.62 (0.77)	9.75 (1.99)	0.36
0.4	11.07 (2.00)	11.07 (2.00)	10.99 (1.93)	1.00
0.6	10.83 (2.79)	10.16 (2.45)	9.45 (1.41)	1.49
0.8	10.55 (3.44)	8.88 (2.45)	8.09 (1.07)	1.91
1.0	9.80 (3.68)	7.66 (2.17)	7.14 (0.88)	2.30
1.2	9.43 (4.14)	6.97 (2.16)	6.48 (0.79)	2.65

Table 5: Idea 3: spike-band noise after sampling, μV (aliasing penalty in dB).

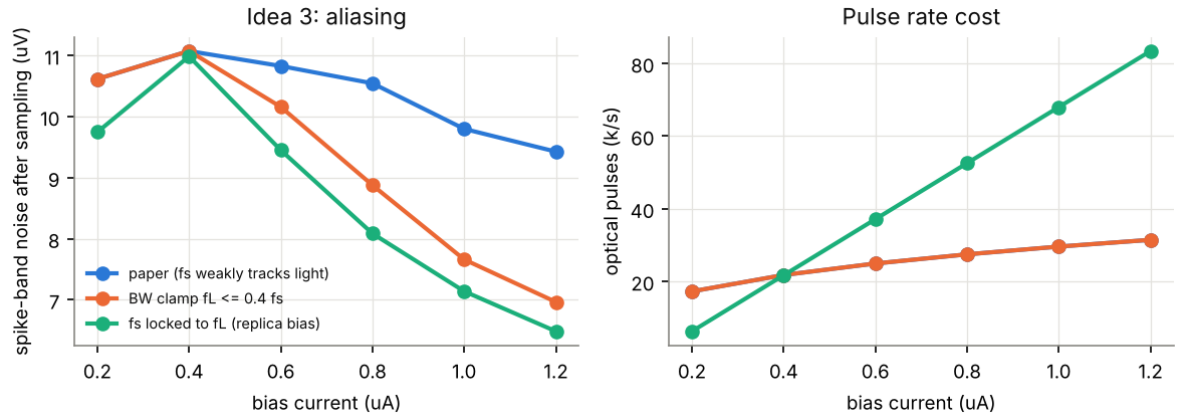

Figure 6: Idea 3: spike-band noise after sampling and optical pulse rate.

4.4 Idea 4: a reference pulse only every K periods

Method. A divide-by-K counter gates the reference pulse and the receiver interpolates the grid with a second-order loop, so pulses per sample fall from 2 to $1 + 1/K$. A SKY130 synchronous Johnson divide-by-8 works from 0.317 to 0.43 V with no glitches at 0.35-0.50 nA; a ripple counter also divides by 8 but produces decode glitches of up to 4.5 μs exactly at the clock edge, where the reference pulse fires. *Result* (Fig. 7, Table 6): $K = 8$ keeps the link error at 0.69 μV (median) with F1 0.837 (K=1: 0.828), 44 % fewer LED pulses. $K = 2$ and $K = 16$ fail (Section 5).

cycle jitter	K	pulses/sample	median link error (μV)	F1
0.0001	1	2.000	0.11	0.828
0.0001	2	1.500	805.00	0.000
0.0001	4	1.250	0.24	0.828
0.0001	8	1.125	0.69	0.837
0.0001	16	1.062	1042.93	0.000
0.001	1	2.000	0.15	0.803
0.001	2	1.500	804.32	0.000
0.001	4	1.250	1.54	0.828
0.001	8	1.125	2.42	0.810
0.001	16	1.062	1007.27	0.000

Table 6: Idea 4 under light flicker, 10 photons per pulse.

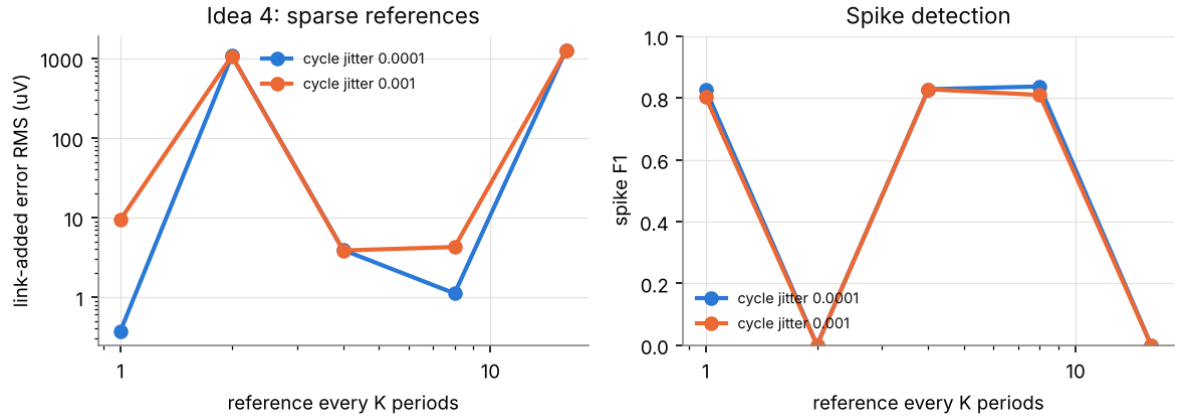

Figure 7: Idea 4: link error and spike F1 versus reference decimation K.

4.5 Ideas 5 and 7: send-on-delta and spike-gated encoding

Method. Send-on-delta [10, 9] sends an encoding pulse only when the signal has moved by Δ; spike-gating samples at 1 kHz and switches to full rate for 2-3 ms when an on-mote high-pass threshold detector fires. Both are evaluated on a perfect link against uniform sampling at the same average rate. *Result* (Fig. 8): with $\Delta = 20\,\mu$V send-on-delta uses 4663 instead of 14798 samples/s at F1 0.855 (full rate 0.847; uniform at the same rate 0.700), a 78 % pulse saving together with Idea 4. Spike-gating saves 89 % of pulses at F1 0.617 (uniform at the same rate 0.292), but loses LFP fidelity (21.3 versus 32.8 dB).

4.6 Idea 6: the oscillator period as a free ripple sensor

Method. The period follows the supply ($d \ln T/dV = -5.35$ /V from Fig. 11), so the receiver already has a reference for supply ripple. I regress the decoded signal on the detrended log-period (8 lags, 0.5 s blocks) and subtract, a reference-based noise canceller [8]. Assumptions: input-referred supply sensitivity 0.01, ripple band 1-1000 Hz; the share ρ of the optical excess noise that is ripple is unknown and swept. *Result* (Fig. 9): $\rho = 1$: IRN 16.09 to 12.74 μV (electrical chain 11.83), F1 0.784 to 0.840; $\rho = 0.5$: 15.37 to 14.72; $\rho = 0$: no gain.

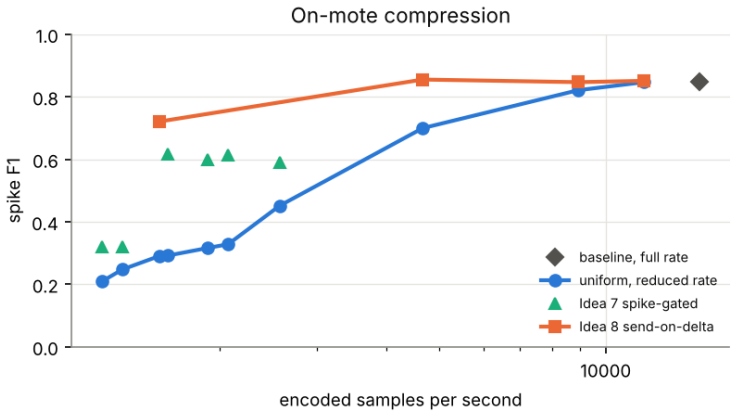

Figure 8: Ideas 5 and 7: spike F1 versus encoded samples per second.

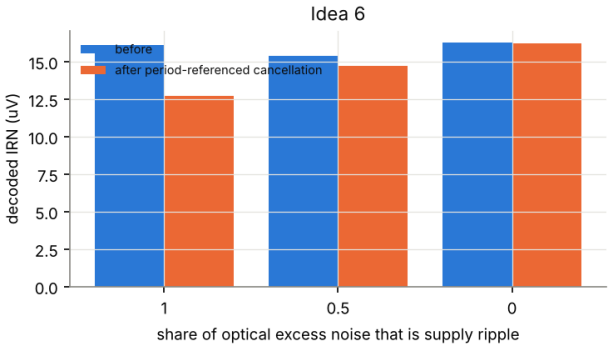

Figure 9: Idea 6: decoded IRN before and after period-referenced cancellation.

4.7 Idea 8: several motes, decoded by their periods

Method. Each mote's period acts as an ID. The receiver finds every reference grid, tracks each with its own loop, excludes pulses claimed as other motes' references and picks encoding pulses by continuity rather than brightness (equal amplitudes, the hard case of Fig. 21 in [1]). *Result* (Table 7, Fig. 10): two motes are both decoded (median F1 0.702 versus 0.053 for a single-node decoder); four or more are not.

motes	grids found	median F1	perfect link	single-node decoder
1	1	0.805	0.805	0.805
2	2	0.702	0.820	0.053
4	4	0.278	0.821	0.000
8	8	0.132	0.803	0.000
16	16	0.059	0.812	0.000

Table 7: Idea 8, equal amplitudes, 20 photons per pulse.

5 Negative results

- **Chopper-first input (Idea 9).** Commutating the electrodes before the amplifier, synchronous with the encoder's pseudo-chopping, by analogy with mixer-first receivers [13], helps only with ideal switches at high light (input-referred LFP noise 5.18 to 3.84 μV, Fig. 11). SKY130 transmission gates at 0.317-0.43 V have 1.13-3.14 MΩ at $W = 10\,\mu$m, whose thermal noise makes it worse (8.24 μV), and at 0.3 μA the amplifier bandwidth is below the chopping frequency and the gain collapses to 0.23 [11, 12].
- **Reference decimation limits.** $K = 2$ fails at acquisition (the same-parity encoding train is indistinguishable from a reference train at $2T$); a distinct reference marker would be needed. $K = 16$ loses lock under the flicker used here.

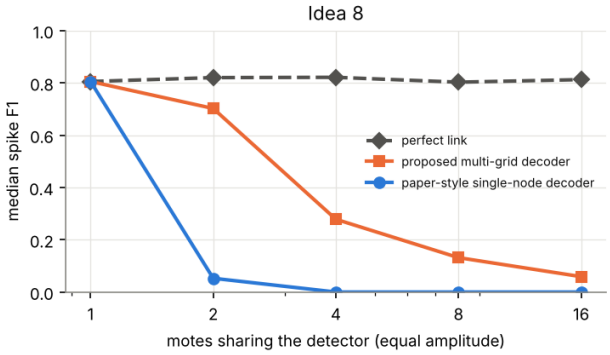

Figure 10: Idea 8: spike F1 versus number of motes sharing one detector.

- **Ratio-locked sampling at low light** samples too slowly for spikes (F1 0.107).
- **Multi-mote decoding** fails for four or more equal-amplitude motes without joint assignment.
- **Ripple cancellation** gives nothing if the excess noise is not supply-coupled ($\rho = 0$; F1 0.774 to 0.764).
- **Spike-gating** keeps fewer spikes than full rate and degrades the LFP.

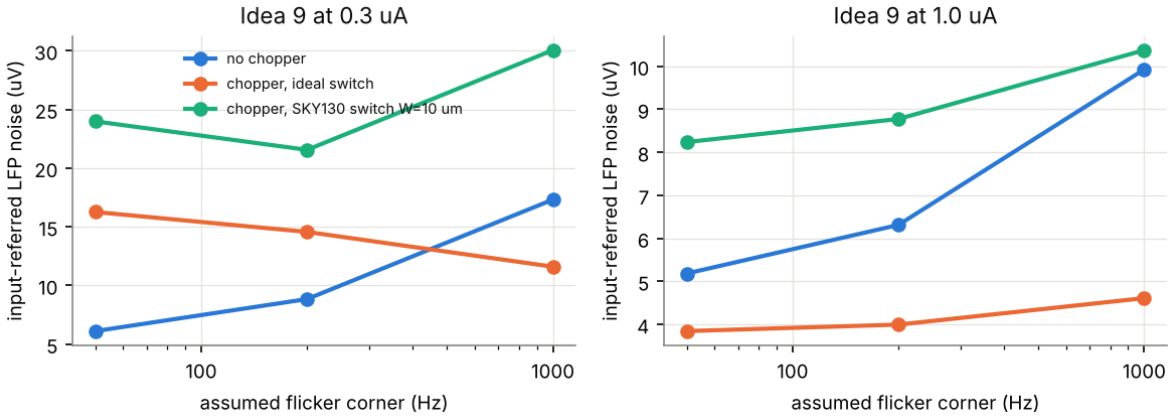

Figure 11: Idea 9 (negative): input-referred LFP noise with and without a chopper, versus the assumed flicker corner.

6 Outlook

The rebuilt chain suggests that the mote is already close to its noise limit and that the next gains are on the link and encoding side: a receiver that tracks the reference grid and uses every photon, a sampling rate kept above twice the bandwidth, and encoding that spends photons only when the signal changes. The obvious next steps are to run the decoders on the group's recorded oscilloscope traces, to measure whether the optical excess noise correlates with the decoded period, and to design the send-on-delta memory and the replica bias at transistor level under illumination. Limits of this study: the photon-counting channel, the 50 Hz flicker corner, the full-scale range and the supply sensitivity are assumptions; the transistor-level parts are SKY130 replicas, not the TSMC 180 nm design.

References

[1] R. Jiao, Y. Zheng, A. M. Nibras, S. Ghajari, S. Sadeghi, A. J. Cortese, P. L. McEuen, A. C. Molnar, and S. Lee, "A 43 μm $\times$ 269 μm light-adaptive optoelectronic autonomous microsystem for neural recording," *IEEE Trans. Biomed. Circuits Syst.*, vol. 20, no. 2, pp. 180-193, Apr. 2026.

[2] S. Lee, A. J. Cortese, P. Trexel, E. R. Agger, P. L. McEuen, and A. C. Molnar, "A 330 μm $\times$ 90 μm opto-electronically integrated wireless system-on-chip for recording of neural activities," in *Proc. IEEE ISSCC*, 2018, pp. 292-294.

[3] S. Lee, A. J. Cortese, A. P. Gandhi, E. R. Agger, P. L. McEuen, and A. C. Molnar, "A 250 μm $\times$ 57 μm microscale opto-electronically transduced electrodes (MOTEs) for neural recording," *IEEE Trans. Biomed. Circuits Syst.*, vol. 12, no. 6, pp. 1256-1266, Dec. 2018.

[4] S. Lee *et al.*, "A subnanolitre tetherless optoelectronic microsystem for chronic neural recording in awake mice," *Nature Electron.*, vol. 8, pp. 1259-1271, Dec. 2025.

[5] H. Hemmati, A. Biswas, and I. B. Djordjevic, "Deep-space optical communications: future perspectives and applications," *Proc. IEEE*, vol. 99, no. 11, pp. 2020-2039, Nov. 2011.

[6] B. Moision and J. Hamkins, "Deep-space optical communications downlink budget: modulation and coding," *IPN Progress Report* 42-154, JPL, Aug. 2003.

[7] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Trans. Syst., Man, Cybern.*, vol. 9, no. 1, pp. 62-66, Jan. 1979.

[8] B. Widrow *et al.*, "Adaptive noise cancelling: principles and applications," *Proc. IEEE*, vol. 63, no. 12, pp. 1692-1716, Dec. 1975.

[9] Y. Tsividis, "Event-driven data acquisition and digital signal processing: a tutorial," *IEEE Trans. Circuits Syst. II*, vol. 57, no. 8, pp. 577-581, Aug. 2010.

[10] M. Miskowicz, "Send-on-delta concept: an event-based data reporting strategy," *Sensors*, vol. 6, no. 1, pp. 49-63, 2006.

[11] 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, pp. 1584-1614, Nov. 1996.

[12] T. Denison *et al.*, "A 2 μW 100 nV/$\sqrt{\text{Hz}}$ chopper-stabilized instrumentation amplifier for chronic measurement of neural field potentials," *IEEE J. Solid-State Circuits*, vol. 42, no. 12, pp. 2934-2945, Dec. 2007.

[13] C. Andrews and A. C. Molnar, "A passive mixer-first receiver with digitally controlled and widely tunable RF interface," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2696-2708, Dec. 2010.

[14] SkyWater Technology and Google, "SKY130 open-source process design kit," 2020. [Online]. Available: github.com/google/skywater-pdk