A VCO-Based Neural Front-End for Body-Coupled Communication: Rebuilt Block by Block, Validated, and Extended (Simulation Study)

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

November 2026

Abstract

This report rebuilds the intravascular neural front-end of Meimandi *et al.* [1] (EPFL Bio/CMOS Interfaces Laboratory with Politecnico di Torino) as a behavioural model, block by block, and ports its LNA and VCO to SKY130 at transistor level. The model reproduces the measured output swing for a $50\,\mu V_{pp}$ tone (237.0 to 254.3 kHz against 237 to 254 kHz), the spike-band noise (5.51 against $5.55\,\mu V_{rms}$) and the dynamic range of Eq. (1) (42.16 against 42.1 dB). Ten ideas are then simulated against this baseline. The strongest result is at transistor level: in a SKY130 replica of the VCO stage, a resistor-degenerated voltage-to-current control reduces the tuning nonlinearity from 30.5% to 1.18% and the spread of the centre frequency over five process corners from 27.3× to 1.35×. Everything here is simulation; the paper's numbers are measured silicon, mine are not.

1 Introduction

Intravascular recording reaches deep brain sources through the blood vessels without opening the skull [2]. A recorder at the tip of a microcatheter must be extremely narrow and must not depend on a signal wire. The letter [1] answers this with an LNA, a current-starved ring VCO and an inverter driver: the neural signal is encoded in the output frequency, $f_{out} = f_0 + K_{vf}\, v_{in}$, so the attenuation α of the body channel scales the carrier amplitude but not the recovered signal (Sec. V-E of [1]). The chip occupies $0.0025\,mm^2$ with a $9.92\,\mu m$ core width and uses $11\,\mu W$ from ± 0.6 V rails delivered through the catheter.

My goals were (i) to understand every block well enough to rebuild it, (ii) to check that the rebuilt model reproduces the published measurements, and (iii) to use it to test new ideas, including circuit changes at transistor level. All code, data and the numbers below are in the accompanying folder; every number in this report comes from a results file.

2 The paper rebuilt

Signal chain. The model follows Fig. 1(b) of [1] in signal order.

1. **Neural signal.** Three units (80, 150 and $300\,\mu V_{pp}$, within the 40 to $500\,\mu V$ range of Sec. I), a local field potential ($40\,\mu V_{rms}$, 1 to 300 Hz) and background activity ($3\,\mu V_{rms}$, 300 Hz to 6 kHz).
2. **LNA.** Input-referred noise $S(f) = S_w + K_f/f^{\alpha}$; band-pass with a shelf high-pass and a low-pass pole; soft output saturation $z = X_s\, u/(1+|u|^p)^{1/p}$, $u = y/X_s$. The T-network gives $C_{f,eff} = C_1 C_2/(C_1 + C_2 + C_{sh})$, i.e. an ideal gain of 1570 V/V if the 35.6 fF shunt sits at a differential virtual ground (920 V/V if grounded); the letter does not state the LNA gain.
3. **VCO.** Logistic tuning $f_v(u) = f_{max}/(1 + R\, e^{-u/U})$ with $f_v(0) = f_0 = 245.5$ kHz and $f_v'(0) = K_0 = 347\,Hz/\mu V$: exponential at low current (the push-pull control branch is in weak inversion, as my SKY130 port shows) and saturating at f_{max}. A white frequency-noise floor after the LNA represents the VCO and the frequency extraction. The VCO is integrated edge by edge, so every edge time is exact and anything clocked by the VCO's own cycle count switches exactly at a cycle boundary.
4. **Body channel.** Not described in [1]; my assumptions: the ± 0.6 V square wave, a first-order electrode high-pass at 100 kHz, a flat loss (swept), receiver noise $10\,nV/\sqrt{Hz}$, a 60 to 600 kHz receiver band-pass.
5. **Receivers.** Exact phase (ideal), an edge counter with $sinc^2$ decimation, comparator plus fine time-to-digital converter, I/Q discriminator, and a type-2 PLL.
6. **Spike detection.** Band-pass 300 Hz to 6 kHz, threshold -4.5σ (median absolute deviation), F1 score against the known spike times.

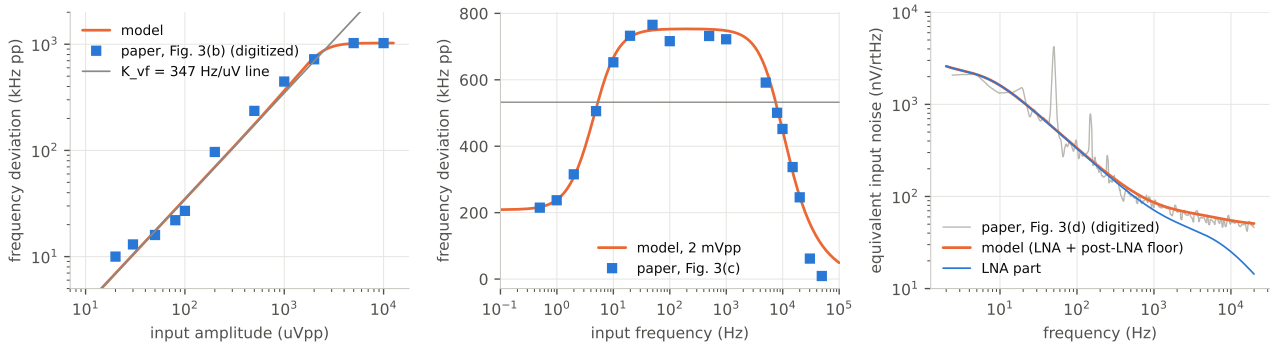


Figure 1: Fit to the digitized Fig. 3 of [1]. Left: peak-to-peak deviation against input amplitude. Middle: deviation against input frequency at $2\,\mathrm{mV_{pp}}$. Right: equivalent input noise, with the LNA part shown separately.

The equations that matter. The LNA's mid-band gain and high-pass corner are

$$A_{\mathrm{mid}} = \frac{C_{\mathrm{in}}}{C_{f,\mathrm{eff}}}, \qquad C_{f,\mathrm{eff}} = \frac{C_1 C_2}{C_1 + C_2 + C_{\mathrm{sh}}}, \qquad f_{\mathrm{HP}} = \frac{1}{2\pi R_{\mathrm{pseudo}} C_{f,\mathrm{eff}}}, \tag{1}$$

with $C_1 = C_2 = 7.4\,\mathrm{fF}$ and $C_{\mathrm{sh}} = 71.2\,\mathrm{fF}$ per half for differential signals. A current-starved ring with N stages, node capacitance C and swing V runs at

$$f = \frac{I_{\mathrm{ring}}}{NCV}, \qquad K_{\mathrm{VCO}} = \frac{\partial f}{\partial V_c} = \frac{1}{NCV}\frac{\partial I_{\mathrm{ring}}}{\partial V_c}, \tag{2}$$

so the tuning law is the voltage-to-current law of the control branch. In weak inversion $I \propto e^{V_{GS}/(nU_T)}$, which makes f_0 and K_{VCO} exponential in the threshold voltage. The noise efficiency factor

$$\mathrm{NEF} = V_{\mathrm{rms,in}}\sqrt{\frac{2I_{\mathrm{tot}}}{\pi U_T\, 4kT\,\mathrm{BW}}} \tag{3}$$

is not reported; with the whole-chain current ($11\,\mu\mathrm{W}/1.2\,\mathrm{V}$) and $5.55\,\mu\mathrm{V_{rms}}$ over 5.7 kHz it is about 8.4 (my estimate). At the receiver, $v_{\mathrm{in}} = (f - f_0)/K_{vf}$; a non-reset edge counter read at f_s and sinc^2-decimated adds a quantisation noise power of about $(2\pi^2/9)B^3 f_0/f_s^2$ (Hz^2) in a band B. Below a carrier-to-noise ratio of about 10 dB in the receiver band, zero crossings are missed and FM demodulation collapses (the FM threshold).

Fitting. I digitized Fig. 3(a) to (d) automatically (axis ticks and square markers detected on a 600 dpi render). With K_0 fixed at $347\,\mathrm{Hz}/\mu\mathrm{V}$, the static curve matches the small-signal point, the $2\,\mathrm{mV_{pp}}$ point (+3.7%) and the saturation, but sits 18 to 28% below the 200 to 1000 $\mu\mathrm{V_{pp}}$ points. A fit with K_0 free matches them to 1.5% but needs $K_0 = 495\,\mathrm{Hz}/\mu\mathrm{V}$, which contradicts Fig. 3(a) ($332\,\mathrm{Hz}/\mu\mathrm{V}$ at $50\,\mu\mathrm{V_{pp}}$). I kept $K_0 = 347\,\mathrm{Hz}/\mu\mathrm{V}$. The band-pass fits Fig. 3(c) to 3.1% up to 20 kHz (shelf high-pass 5.95 Hz, low-pass 7.37 kHz). The noise fit splits the $5.55\,\mu\mathrm{V_{rms}}$ into $4.15\,\mu\mathrm{V_{rms}}$ from the LNA and $3.68\,\mu\mathrm{V_{rms}}$ from a flat floor of about $49\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$, which Fig. 3(d) shows up to 20 kHz, above the 6.7 kHz bandwidth (Fig. 1).

3 Baseline validation

Table 1 compares the time-domain model with the paper. The minimum detectable signal of $20\,\mu\mathrm{V_{pp}}$ corresponds to a spike-band SNR of 2.4 dB in the model, which is 0 dB at $15.3\,\mu\mathrm{V_{pp}}$. The wide-band noise is $9.85\,\mu\mathrm{V_{rms}}$ without mains and $13.88\,\mu\mathrm{V_{rms}}$ with the 50, 150 and 250 Hz lines read from Fig. 3(d) (15.5 measured). Figure 2 overlays the model on the four panels of Fig. 3.

Link budget. With the assumptions above, frequency encoding keeps the ideal-receiver noise up to 85 dB of channel loss (carrier-to-noise ratio 11.7 dB) and collapses between 88 and 91 dB. An amplitude-domain link with $A_v = 1000$ degrades from 70 dB (Fig. 3). This quantifies the argument of Sec. V-E: frequency encoding buys about 15 dB here, and its limit is the FM threshold.

SKY130 transistor-level ports of the paper's blocks

Table 2 lists the substitutions; every width and length is from Table 1 of [1] unless stated. I ported Fig. 2(a) to (c) with the Table 1 sizes at ±0.6 V, using SKY130 low-threshold 1.8 V devices, PMOS lengths of at least

Quantity	Paper (measured)	Model (simulated)	In [1]
Output frequency, 50 μV_{pp}, 1 kHz	237 to 254 kHz	237.0 to 254.3 kHz	Fig. 3(a)
Spike-band noise (300 Hz to 6 kHz)	5.55 μV_{rms}	5.51 μV_{rms}	Sec. V-D
Minimum detectable signal	20 μV_{pp}	SNR 2.4 dB at 20 μV_{pp}	Sec. V-B
Bandwidth (2 mV_{pp})	6 Hz to 6.7 kHz	5.4 Hz to 7.5 kHz	Fig. 3(c)
Linear range	2 mV_{pp}	1-dB compression at 3.08 mV_{pp}	Fig. 3(b)
Dynamic range, Eq. (1)	42.1 dB, 6.7 bits	42.16 dB, 6.71 bits	Sec. V-E
Spike detection F1	not reported	0.985	

Table 1: Baseline: the rebuilt chain against the measured numbers.

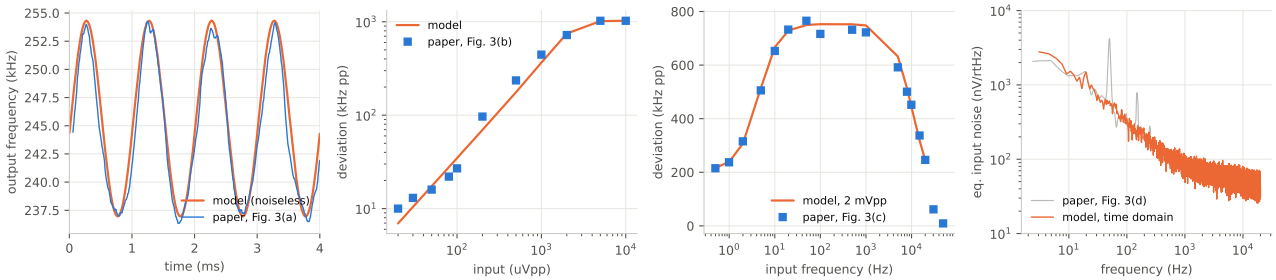


Figure 2: Time-domain model against Fig. 3(a) to (d) of [1].

0.35 μm, standard-threshold pseudo-resistors, 20 MΩ resistors for the common-mode sensing devices M5/M6, and ideal capacitors. With M0's gate at ground as drawn, the LNA port gives 56.7 dB (683 V/V) from 45.8 Hz to 10.6 kHz, 2.51 μV_{rms} of its own noise in the spike band, and 13.5% THD at 2 mV_{pp}, consistent with the paper's statement that the LNA saturates above 2 mV_{pp}. The faithful VCO port runs at 7.14 kHz at zero input instead of 245 kHz, because the series push-pull branch works in deep weak inversion and SKY130 thresholds are higher than TSMC 180 nm's. Its frequency changes exponentially with input and threshold: 1.39 to 38.0 kHz over five corners and 7.05 %/°C. The ring alone, fed by an ideal current, is close to linear: 3.09 to 3.41 kHz/nA from 10 to 320 nA. So the nonlinearity and the threshold sensitivity come from the voltage-to-current step, not from the ring.

4 Ideas

The ideas are ranked from safe to ambitious. Each was simulated against the baseline above.

Idea 1: receiver look-up-table inverse

Why. The implant is the scarce resource and the receiver is not; if the static transfer curve is known, the receiver can undo it without touching the chip. *Method.* 1 kHz calibration sines at eight amplitudes up to 5 mV_{pp} build a 300-bin table from frequency deviation back to the band-passed input, which the receiver applies to every sample. *Result.* At 2 mV_{pp} the THD falls from 26.3% to 0.38%, and the peak SNDR rises from 21.7 to 35.0 dB. Above 3 mV_{pp} the inverse amplifies noise where the curve saturates. The baseline THD depends on

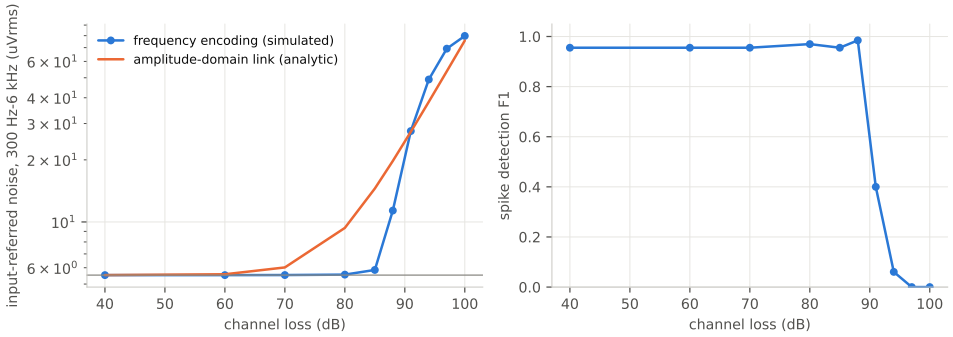


Figure 3: Link budget over the body channel: input-referred noise (frequency encoding simulated, amplitude-domain link analytic) and spike F1 against channel loss, with the assumptions of Sec. 2.

Element	Paper (TSMC 180 nm)	SKY130 port
All MOSFETs	1.8 V devices, Table 1	low-V_t `nfet_01v8_lvt`, `pfet_01v8_lvt`
M0, M43, M44 (L = 0.25, 0.18 μm)	PMOS	L = 0.35 μm (minimum for the low-V_t PMOS)
Pseudo-resistors M13 to M20 (2/16)	PMOS diodes	standard-V_t `pfet_01v8` (lower leakage)
M5/M6 common-mode sensing (1/8)	MOS resistors	20 MΩ resistors (drawn version floats at DC)
Output driver M46 to M49	medium-V_t, size not given	2/0.35 and 1/0.15, then 6/0.35 and 3/0.15
Capacitors	MIM / MOM	ideal

Table 2: SKY130 port of Fig. 2: substitutions.

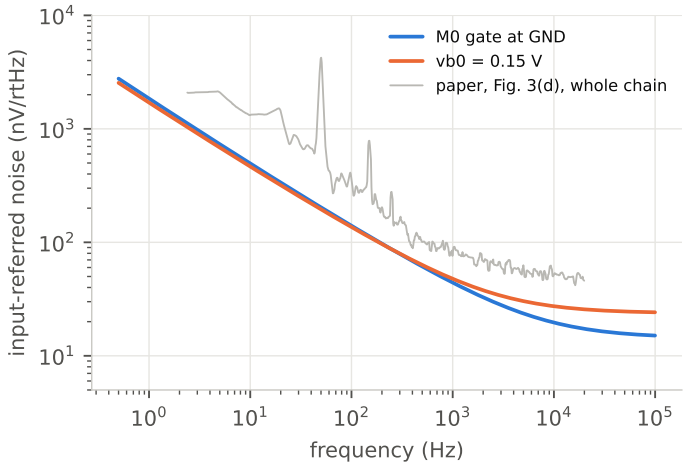


Figure 4: Input-referred noise of the SKY130 LNA port (two biases) against the whole-chain equivalent input noise of Fig. 3(d) of [1]. The port's own noise lies below the measured chain noise, consistent with a post-LNA contribution to the measured floor.

my exponential-VCO assumption: peak-to-peak data cannot show even-order distortion, and one spectrum measurement would settle it. *Risk.* The table must be refreshed as the curve drifts with temperature; Idea 5 can provide the calibration signal in vivo.

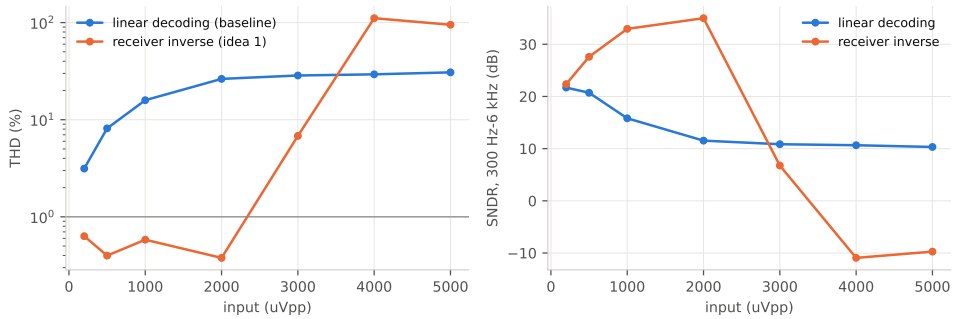


Figure 5: Idea 1: THD and SNDR of a 1 kHz tone with linear decoding and with the receiver look-up-table inverse.

Idea 2: counter receiver

Method. A non-reset edge counter read at f_s, followed by sinc^2 decimation to 50 kS/s. The predicted quantisation noise is $(2\pi^2/9)B^3 f_0/f_s^2$. *Result.* The counter adds 1.36 μV_{rms} at 1 MHz and nothing resolvable from 2 MHz up (0.46 μV_{rms} predicted at 2 MHz). Spike F1 is 1.000 at every clock. A 2 to 10 MHz counter is enough for this link. The idea is modest on its own; it defines the external unit and is the receiver assumed by Ideas 4 and 7.

Idea 3: resistor-defined voltage-to-current control (transistor level)

Method. The load-balancing amplifier and M27/M28 are replaced by a source-degenerated NMOS pair driven by V_{o+}/V_{o-}, with a resistor tail. Its V_{o+}-side current, $I_a = I_t/2 + v_d/(R_s + 2/g_m)$, is the ring's starving current through the existing M29/M26 mirrors. The ring and buffers are unchanged. Resistors are SKY130 xhigh-poly:

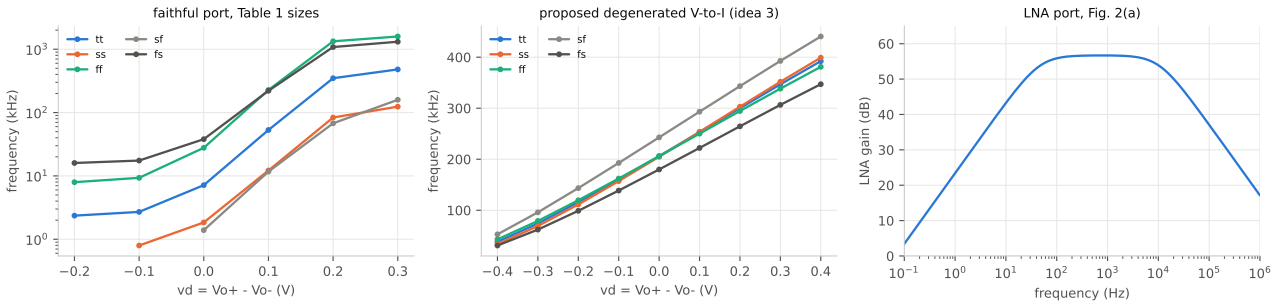


Figure 6: SKY130 ports. Left: faithful VCO port (Table 1 sizes), frequency against differential LNA output over five corners (log scale). Middle: proposed degenerated driver with the same ring (linear scale). Right: LNA port gain with M0's gate at ground.

2.8 MΩ per side and 1.43 MΩ tail, about 430 μm^2 drawn. *Result* (Table 3, Fig. 6): the nonlinearity over ±0.3 V falls from 30.5% to 1.18%. The f_0 spread over five corners falls from 27.3× to 1.35×, the K_{VCO} spread from 19.1× to 1.20×, and the f_0 temperature coefficient from 7.05 to 0.22 %/°C, at 1.42 instead of 1.73 μA. *Why it works.* The degenerated pair converts voltage to current with $G_m \approx 1/(R_s + 2/g_m)$, which depends on a resistor rather than on an exponential in V_{th}, and the resistor tail sets I_t linearly in V_{th}. The ring then sees a well-defined current, and the ring itself is linear in current. *Risk.* The resistors cost about a sixth of the paper's 2502 μm^2 (they fit the 10 μm width when folded), mismatch and noise of the degenerated pair are not yet simulated, and the TSMC 180 nm result may differ. On the paper's chip the measured linear range is set by the LNA, so the gain is mainly robustness.

SKY130, ±0.6 V, same ring	Table-1 port	Proposed
f_0 at tt, 27°C	7.14 kHz	206.1 kHz
f_0 spread over tt/ss/ff/sf/fs (max/min)	27.3×	1.35×
K_{VCO} spread over the five corners	19.1×	1.20×
f_0 temperature coefficient, 27 to 47°C	7.05 %/°C	0.22 %/°C
Nonlinearity, $v_d = \pm 0.3$ V (% of span)	30.5%	1.18%
Supply current at $v_d = 0$	1.73 μA	1.42 μA

Table 3: Idea 3 at transistor level (schematic simulation, ngspice, SKY130 tt/ss/ff/sf/fs).

Idea 4: move gain from the LNA into the frequency domain

Method. Divide the LNA gain by G, which stretches the transfer curve by G, so post-LNA noise is referred to the input G times larger. *Result* (Fig. 7, right): this is conditional on the origin of the flat floor. If 1% of it remains (it was extraction), $G = 10$ gives a 30.9 mV$_{\mathrm{pp}}$ range at 5.52 μV$_{\mathrm{rms}}$ and 65.9 dB, against 45.9 dB at $G = 1$. If all of it remains (it is VCO noise), the noise rises to 36.8 μV$_{\mathrm{rms}}$ at $G = 10$, a negative result for spikes. The first measurement to make is therefore the noise floor above 7 kHz with a fine counter receiver: if it drops, it was extraction and this idea is worth about 20 dB of dynamic range.

Idea 5: carrier-locked test tone

Method. The VCO output divided by 256 gives a 0.959 kHz square wave. Mode A injects it at the LNA input (±25 μV) to measure K_{vf}. Mode B drives the electrode through 7.4 fF (electrode model 10 MΩ ∥ 150 pF) to measure $|Z|$. The receiver knows the tone's phase from the cycle count and runs a lock-in in the cycle domain. *Result* (Fig. 8, left): after a −20% drift, K_{vf} is recovered within 2.83% (0.1 s), 1.19% (1 s) and 0.24% (3 s) with the neural signal present, and $|Z|$ within 0.33 to 0.72%. Spike F1 is 0.989 after the tone is subtracted, against 0.208 if it is left in. *Why it is new.* Pilot tones and impedance checks exist in neural SoCs; deriving the tone from the transmitted carrier makes the external receiver synchronous for free, which matters for a node with no clock, no digital back end and no wire. It fits the authors' electrode modelling [4]. *Risk.* It needs a femtofarad-scale injection network and a tone-off window for the most sensitive recordings.

Idea 6: self-clocked chopping at the VCO input

Method. V_{c+} and V_{c-} are swapped every four VCO cycles (31 kHz). The receiver splits the cycles by chopper state, interpolates both streams on one grid and takes $(f_+ - f_-)/2$, which keeps only odd orders. *Result:* at

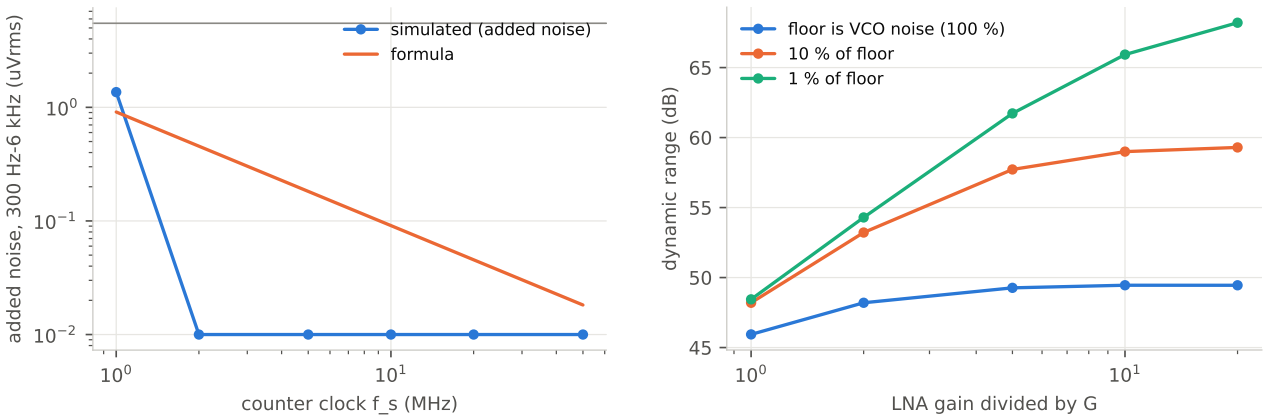


Figure 7: Left: Idea 2, noise added by the counter receiver against its clock, simulated and predicted. Right: Idea 4, dynamic range against the LNA gain reduction for three assumptions on the noise floor.

$1\,\mathrm{mV_{pp}}$ the HD2 falls from 15.9% to 0.20% and the THD to 1.2% (third-order limited). The in-band residual of a 0.2% VCO drift at 2 Hz falls from 0.99 to about $4 \times 10^{-6}\,\mu\mathrm{V_{rms}}$. My first receiver weighted each cycle by its duration and turned HD2 into HD3; separating the two streams fixed it. *Risk.* The transmitted spectrum widens by the chopping rate, and errors before the chopper (the LNA's) are not removed; the cost on the implant is four switches and two flip-flops.

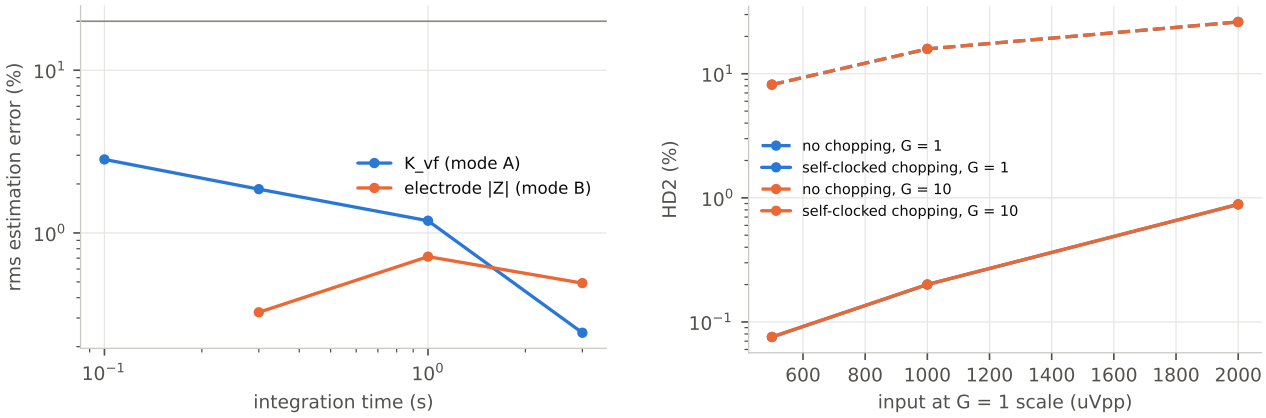


Figure 8: Left: Idea 5, estimation error of K_{vf} and of the electrode impedance against integration time. Right: Idea 6, second harmonic with and without self-clocked chopping, at the paper's gain and with the LNA gain divided by 10.

Idea 7: PLL receiver

Result (Fig. 9): the noise stays within 1 dB of ideal up to 94 dB of loss with a 20 kHz type-2 PLL, against 91 dB for an I/Q discriminator and 85 dB for the comparator receiver. That is 9 dB of extra margin. The absolute losses depend on my channel assumptions; the difference between receivers is the robust result. The PLL works because it only admits noise inside its loop bandwidth, which is matched to the 6.7 kHz signal rather than to the wide FM spectrum.

Idea 8: frequency-division multi-channel link

Four chains with $K = 87\,\mathrm{Hz}/\mu\mathrm{V}$ on carriers within one octave share one receiver. For a $1\,\mathrm{mV_{pp}}$ tone the worst crosstalk is -44 dB at 40 kHz spacing, -25 dB at 60 kHz and -57 dB at 90 kHz. The spike F1 is 0.68 to 0.94 per channel, against 0.64 for a single channel at the same reduced gain. Crosstalk is therefore not the limit; the gain reduction with the paper's noise floor is (partly negative). Ref. [5] already sends multisite data over galvanic body-channel communication; the novelty here is weak until its access scheme is checked.

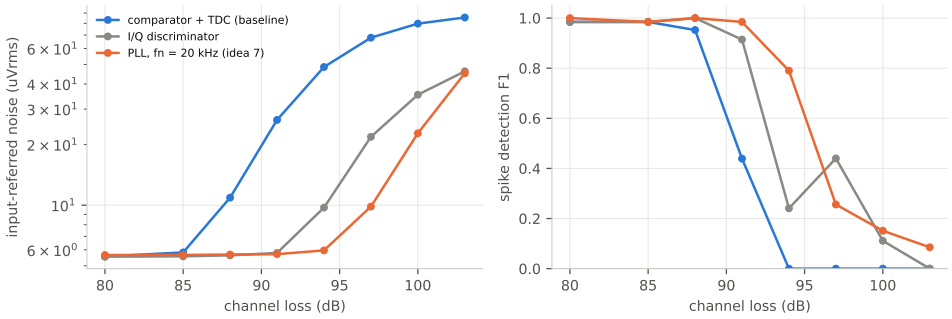


Figure 9: Idea 7: three receivers on the same received waveforms; input-referred noise and spike F1 against channel loss.

Idea 9: keyed frequency mask

A 6-bit DAC adds a PUF-keyed offset every 16 cycles (0.5 LSB INL). The key holder learns each code's offset and subtracts it. With a $\pm 50\,\mu$V-equivalent mask, F1 is 0.976 for the key holder, 0.316 for an eavesdropper and 0.209 for a level-grid attack. At $\pm 400\,\mu$V the mask uses up VCO range and the key holder drops to 0.420 (negative). A secure keystream generator may cost more area than the whole chip, and masking is weaker than digital encryption; the idea connects to my PUF-keyed in-modulator encryption for delta-sigma modulators.

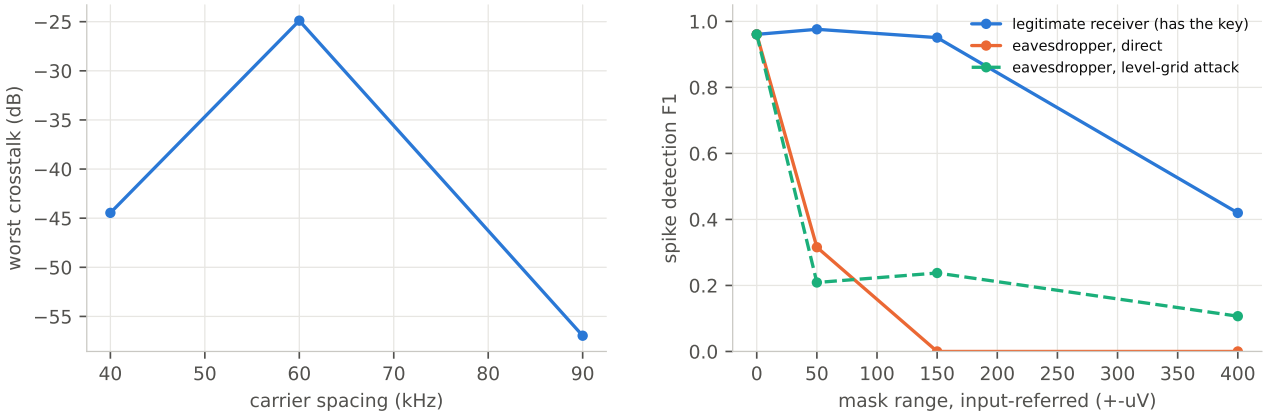


Figure 10: Left: Idea 8, worst crosstalk against carrier spacing. Right: Idea 9, spike F1 for the key holder and for an eavesdropper against the mask range.

Idea 10: the carrier as a position beacon

In a calibration mode, 32 electrodes measure the carrier amplitude and a current-dipole fit locates the tip. This is a homogeneous-medium model, so the results are an optimistic bound. The median error is 5.5 mm after 1 s with the on-chip 0.2 mm electrode spacing, 1.7 mm after 10 s, and 0.55 mm after 1 s with 2 mm spacing. The skull and the vessel geometry are not modelled, so the next step is a three-shell head model; clinically the value would be monitoring migration rather than replacing fluoroscopy.

5 Negative and conditional results

- The receiver inverse (Idea 1) fails above 3 $\mathrm{mV_{pp}}$, where the curve flattens and noise is amplified.
- Gain re-partition (Idea 4) makes things worse if the flat noise floor is VCO noise: 36.8 $\mu\mathrm{V_{rms}}$ at $G = 10$.
- Frequency division (Idea 8) is limited by the gain reduction, not by crosstalk.
- A large keyed mask (Idea 9, $\pm 400\,\mu$V) defeats the key holder too.
- The faithful SKY130 VCO port cannot reach 245 kHz at ± 0.6 V. Neither scaling the ring mirrors nor raising the supply to ± 0.75 V re-centred it (34.5 kHz), so I compare normalised metrics.
- My model misses the 200 to 1000 $\mu\mathrm{V_{pp}}$ points of Fig. 3(b) by 18 to 28%.

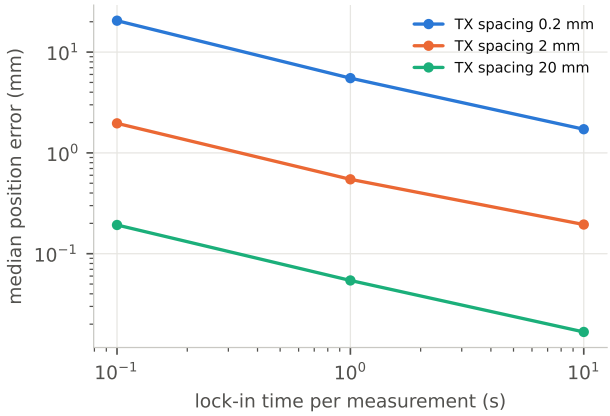


Figure 11: Idea 10: median position error against lock-in time for three transmit electrode spacings (homogeneous medium).

6 Outlook

The ideas fall into three directions that could be validated in silicon. The first is robustness through resistor-defined control (Idea 3) and in-vivo calibration that needs no wire (Idea 5). The second is co-design of the implant and the external receiver (Ideas 1, 2, 4, 6, 7), where resolution and linearity are cheap. The third is scaling the link to many channels, private carriers and position monitoring (Ideas 8 to 10). The first measurements I would make on the existing chip are a spectrum of the demodulated 1 kHz tone (even-order distortion), the noise floor above 7 kHz with a fine counter receiver (noise split), and the catheter-to-receiver channel loss in a phantom.

References

[1] A. Meimandi, P. S. Crovetti, E. Citraro, F. P. Andriulli, and S. Carrara, "A 0.0025 mm^2 CMOS neural front-end with VCO-based frequency encoding for body-coupled communication," *IEEE Sensors Lett.*, vol. 10, no. 9, Art. no. 6008404, 2026, doi: 10.1109/LSENS.2026.3716537.

[2] T. J. Oxley *et al.*, "Minimally invasive endovascular stent-electrode array for high-fidelity, chronic recordings of cortical neural activity," *Nature Biotechnol.*, vol. 34, no. 3, pp. 320-327, 2016.

[3] K. A. Ng and Y. P. Xu, "A compact, low input capacitance neural recording amplifier with C_{in}/gain of 20 fF·V/V," in *Proc. IEEE BioCAS*, 2012, pp. 328-331.

[4] A. Meimandi, G. L. Barbruni, and S. Carrara, "A more reliable equivalent circuit of electrochemical sensors for robust design of CMOS front-ends," *IEEE Sensors J.*, vol. 25, no. 10, pp. 17275-17282, 2025.

[5] D.-H. Choi *et al.*, "An implantable peripheral nerve interface system with simultaneous multisite data transmission via galvanically coupled body-channel communication for next-generation prosthetic limbs," *IEEE Open J. Solid-State Circuits Soc.*, vol. 5, pp. 324-334, 2025.

[6] R. Q. Quiroga, Z. Nadasdy, and Y. Ben-Shaul, "Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering," *Neural Comput.*, vol. 16, no. 8, pp. 1661-1687, 2004.

[7] A. Hajimiri, S. Limotyrakis, and T. H. Lee, "Jitter and phase noise in ring oscillators," *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 790-804, 1999.

[8] SkyWater Technology and Google, "SKY130 open-source process design kit," open_pdks build 1689ac3f, 2026.