

Document Version

Final published version

Licence

Dutch Copyright Act (Article 25fa)

Citation (APA)

Huang, X., Yassin, H., Lafuente Alcazar, B., Akhoundi, A., & Muratore, D. G. (2026). A Bidirectional Neuromodulation Chipset With Algorithm-Aware AFE Optimization and High-Voltage-Compliant Stimulation. *IEEE Journal of Solid-State Circuits*, 61(4), 1538-1550. <https://doi.org/10.1109/JSSC.2025.3650066>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

A Bidirectional Neuromodulation Chipset With Algorithm-Aware AFE Optimization and High-Voltage-Compliant Stimulation

Xiaohua Huang^{ID}, *Member, IEEE*, Hoda Yassin^{ID}, *Student Member, IEEE*, Beatriz Lafuente Alcazar, Arash Akhouni^{ID}, *Graduate Student Member, IEEE*, and Dante G. Muratore^{ID}, *Senior Member, IEEE*

Abstract—This work presents a bidirectional neuromodulation chipset with 64-channel neural analog front-end (AFE), and a four-channel current stimulator. The chipset employs a heterogeneous architecture, combining a 28-nm low-voltage (LV) CMOS process for the AFE and the digital backend (DBE) to improve area and power efficiency, with a 180-nm high-voltage (HV) bipolar-CMOS-DMOS (BCD) process for the stimulation driver to achieve HV compliance. An AFE–DBE co-design framework is proposed to relax AFE noise requirements without sacrificing classification performance in the DBE, thus enabling significant reductions in AFE’s area and power consumption. In addition, we introduce a novel embedded current-steering digital-to-analog converter (IDAC) structure that addresses key limitations of conventional IDACs in multichannel direct-digitization recording architectures, including noise degradation from IDAC switching, and channel-to-channel gain variations. Overall, the proposed chipset achieves HV compliance (10 V) in the stimulator, while maintaining exceptional area and power efficiency in the AFE and DBE. Notably, the AFE achieves the smallest per-channel area (0.0009 mm²) and the lowest per-channel power (204 nW) reported to date for electrocorticography (ECoG) acquisition (≤ 1 kHz bandwidth). Together, these results highlight the chipset’s potential for implantable neuromodulation systems that enable long-term monitoring of neurological disorders with closed-loop intervention.

Index Terms—Analog front-end (AFE), bidirectional, bipolar-CMOS-DMOS (BCD), chipset, closed-loop neuromodulation, CMOS, electrocorticography (ECoG), embedded current-steering digital to analog converter (IDAC), high-voltage (HV) compliance, neurological disorder, stimulation.

I. INTRODUCTION

CLOSED-LOOP neuromodulation systems have emerged as an important therapeutic approach to the treatment of neurological disorders such as epilepsy [1], major depression [2], Parkinson’s disease [3], and memory impairment [4]. These systems typically rely on electroencephalography (EEG) or electrocorticography (ECoG) for continuous monitoring of brain activity, enabling real-time detection of symptoms and

adaptive stimulation [4], which can offer improved treatment efficacy, reduced side effects, and, particularly in implantable and wearable devices, extended battery life [5].

As a result, the past decade has witnessed significant progress in the development of such closed-loop neuromodulation system-on-chips (SoCs), which typically integrate an analog front-end (AFE) for signal conditioning and digitization, followed by a digital back-end (DBE) for symptom detection, and a responsive neurostimulator to close the therapeutic loop [6], [7], [8], [9], [10], [11], [12], [13], [14]. However, existing solutions continue to face critical limitations in the following aspects.

A. Limited AFE Scalability to High Channel Counts

Conventional neural recording pipelines, instrumentation amplifiers followed by analog-to-digital converters (ADCs), are analog intensive and scale poorly with CMOS technologies [15]. Direct digitization approaches improve area and power efficiency, yet their scalability to high channel counts remains challenging with stringent noise and resolution targets [16]. Electrode-level multiplexing reduces per-channel area by sharing a single readout circuitry between several multiplexed electrodes [17], [18], [19], [20], but suffers from limited rejection of electrode dc offsets (EDOs) and reduced input impedance (typically ~ 100 mV_{pp} and tens of M Ω), and does not offer power savings under first-order analysis [19].

B. Poor Channel-to-Channel Uniformity

To improve area and power efficiency for scalability, many recent neural AFEs have adopted open-loop architectures. However, the inherent open-loop amplification introduces significant channel-to-channel variations [14], [17], [21], [22], [23]. Direct-ADC architectures, particularly those based on incremental ADCs, have also gained significant attention for their improved hardware efficiency in scaled CMOS technologies. Nonetheless, they still remain susceptible to poor channel-to-channel uniformity [24], [25]. The resulting large gain errors typically require calibration, adding area and power overhead.

C. Insufficient AFE–DBE Co-Design

Existing AFE designs are typically prioritized for high resolution and low noise [8], [16], [19], [26], aligning with

Received 8 September 2025; revised 2 December 2025; accepted 22 December 2025. Date of publication 12 January 2026; date of current version 30 March 2026. This article was approved by Associate Editor John Wu. This work was supported in part by TU Delft Artificial Intelligence (AI) labs Program and European Union (EU) under Grant 101070931. (*Corresponding author: Xiaohua Huang.*)

The authors are with the Department of Microelectronics, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: X.Huang-6@tudelft.nl).

Digital Object Identifier 10.1109/JSSC.2025.3650066

0018-9200 © 2026 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.

Authorized licensed use limited to: TU Delft Library. Downloaded on April 13, 2026 at 06:20:10 UTC from IEEE Xplore. Restrictions apply.




Fig. 1. Envisioned closed-loop neuromodulation system for the long-term management of neurological disorders, integrating a flexible high-electrode-count μ ECoG array for neural sensing, with stacked CMOS ASICs for signal acquisition, real-time symptom detection, and responsive neurostimulation. A rechargeable microbattery and inductive link provide local power, wireless recharging, and bidirectional data telemetry, all within a biocompatible encapsulation for full implantation.

the requirements of general-purpose neuroscience research. However, such stringent resolution and noise targets may be unnecessary for application-specific tasks such as deducing the brain state. Recent studies have shown that relaxing certain AFE metrics does not significantly degrade decoding performance in the DBE [27], [28], highlighting the importance of tighter AFE–DBE co-design to achieve greater area and power reductions.

D. Limited Voltage Compliance for Stimulation

To benefit from area and power scaling, closed-loop neuromodulation SoCs are increasingly implemented in advanced CMOS technologies [13], [14]. However, the reduced gate oxide thickness significantly lowers the gate breakdown voltage, thereby reducing the available voltage headroom for electrical stimulation. Although transistor stacking can be employed to extend the voltage compliance [14], the maximum achievable value remains limited by the breakdown voltage of the parasitic diodes.

To address the limitations outlined above, Fig. 1 presents our envisioned closed-loop neuromodulation system for the long-term management of chronic neurological disorders. The system integrates a flexible μ ECoG array with compact low-power CMOS ASICs, enabling full implantation and autonomous operation. A high electrode count is essential to achieve both fine spatial resolution and broad cortical coverage [30], enabling more accurate localization of abnormal brain activity and improved therapeutic efficacy. The ASICs employ a heterogeneous integration approach: the AFE and DBE are implemented in scaled low-voltage (LV) CMOS to achieve high area and power efficiency, whereas the stimulation circuitry (STIM) and power management unit (PMU) use a high-voltage (HV) technology to provide the voltage compliance required for effective neural stimulation. The two dies are vertically integrated via 3-D stacking using flip-chip (FC) and through-silicon-via (TSV) bumps, together with local wire bonds, resulting in a compact form factor. A miniature rechargeable microbattery provides local power and is recharged periodically through an inductive coil coupled with




Fig. 2. System-level diagram of the proposed chipset architecture for closed-loop neuromodulation.

a transcutaneous link, while the PMU manages rectification, regulation, and safe charging. The transcutaneous link also supports bidirectional data telemetry for configuration and logging. Overall, the implant performs continuous neural monitoring and real-time symptom detection, and delivers spatially targeted stimulation to suppress abnormal brain activity.

As a proof of concept, Fig. 2 presents our proposed closed-loop neuromodulation chipset architecture for the treatment of epileptic seizures. The chipset integrates a 64-channel AFE for neural signal acquisition, a lightweight DBE for feature extraction and seizure classification, and a four-channel current stimulator for delivering responsive neurostimulation. Each stimulation channel is implemented with a current-steering digital-to-analog converter (IDAC) and a dedicated current driver. The IDAC is co-integrated with the AFE and DBE in a scaled 28-nm LV CMOS technology (Chip A) to improve area and power efficiency, while the current driver is implemented in a 180-nm HV bipolar–CMOS–DMOS (BCD) technology (Chip B), enabling HV compliance with stimulation requirements across a wide range of current amplitudes and electrode–tissue impedances.

It is worth noting that chipset architectures have been reported for various biomedical applications, including neural interfaces; however, prior work either: 1) implement the AFE and DBE on two separate dies within the same CMOS node (i.e., without heterogeneous integration) [12], [31]; or 2) employ heterogeneous CMOS primarily to reduce DBE leakage power by implementing it in an older technology node [32]. In contrast, our proposed heterogeneous partitioning aims to enable high area and power efficiency in the AFE and DBE, while maintaining HV compliance for the stimulation circuitry. Importantly, realizing these benefits requires addressing several system-level challenges, including: 1) increased integration and packaging complexity (e.g., stacking, bump, and interconnect) and 2) interface overhead between two stacked chips such as latency and synchronization constraints.

This work extends [29], with a focus on the design of the recording AFE and the stimulation circuitry. The remainder of this article is organized as follows. Section II presents the




Fig. 3. Left: classifier performance versus added noise, demonstrating robustness up to $\sim 10 \mu\text{V}_{\text{rms}}$ before noticeable degradation. Results are obtained from the software model of [29], and validated using a leave-one-group-out cross-validation and 2-s data windows (nonoverlapping in interictal periods and 90% overlap during ictal segments). Right: channel power versus input-referred noise, comparing this work (target region) with prior neural AFEs with ECoG bandwidth ($\leq 1 \text{ kHz}$).

AFE–DBE co-design. Section III details the AFE. Section IV describes the HV-compliant stimulation. Section V reports the measurement results, and Section VI concludes the article.

II. AFE–DBE CO-DESIGN

In high-channel-count neural recording systems, the AFE’s area and power can grow quickly and thus dominate the overall resource budget. Therefore, techniques for effective area and power reductions are essential. As shown in Fig. 3 (right), existing ECoG AFEs prioritize low-noise recordings [14], [16], [18], [19], [26], [33], [34], achieving noise levels of $\sim 1\text{--}3 \mu\text{V}_{\text{rms}}$ at the expense of high power; designs targeting high dynamic range (DR) suffer from an even less favorable power-noise tradeoff [35], [36], [37]. To achieve such low noise, input chopping or correlated double sampling (CDS) is almost exclusively employed to suppress flicker noise, without which prohibitively large devices would be required. Nevertheless, even with chopping or CDS, transistors must still be sized reasonably large to ensure a sufficiently low flicker-noise corner, incurring substantial area overhead and thus limiting scalability in high-channel-count systems. Furthermore, these techniques reduce the input impedance and provide limited rejection of EDOs.

While these low-noise recordings are valuable for neuroscientific exploration, they are not necessarily essential for certain clinical purposes, such as brain state classification. In fact, recent work has shown that for clinical brain–machine interfaces (BMIs), the AFE noise can be relaxed without degrading DBE decoding performance [27].

Given the stringent noise–power–area tradeoffs, this work adopts an AFE–DBE co-design strategy to prevent over-designing the noise performance and thereby enable reductions in AFE’s power and area. To evaluate this approach, the performance of the neural tree classifier described in [29] is evaluated by adding artificial white noise with a standard deviation of $2\text{--}14 \mu\text{V}_{\text{rms}}$ to the raw iEEG data from the dataset in [38]. For reference, the classifier is a sparse projection oblique randomer forest (SPORF) [39] trained with model-aware feature selection that uses both temporal and synchronization features. As shown in Fig. 3 (left), the classifier’s performance




Fig. 4. Block schematic of the proposed AFE channel and its timing diagram.

is largely unaffected up to $\sim 10 \mu\text{V}_{\text{rms}}$, with noticeable degradation only above this level. A parallel analysis with flicker noise shows comparable robustness over a similar noise range. These results suggest that, by relaxing the noise specification within tolerable limits rather than targeting the state-of-the-art levels of $\sim 2 \mu\text{V}_{\text{rms}}$, significant power savings can be achieved [see the uncharted territory at the bottom-right region of Fig. 3 (right)]. Furthermore, the substantially relaxed flicker noise requirement enables the use of much smaller transistors, and even eliminates the need for input chopping or CDS. As a result, significant area reductions can also be achieved. When input chopping or CDS is not employed, a high input impedance can be preserved, and ac-coupling further provides rail-to-rail EDO rejection.

Classifier robustness to channel-to-channel gain error is also evaluated. Without added noise, the performance remains within 2% of baseline for gain errors up to $\sim 15\%$ (1σ). With added $10 \mu\text{V}_{\text{rms}}$ white noise, the tolerable gain error reduces to $\sim 10\%$ (1σ). Thus, larger gain errors are acceptable in low-noise conditions, whereas higher noise levels require smaller gain errors to maintain the decoding performance.

Finally, it is important to note that while SPORF classifier was employed in this work for its hardware efficiency and interpretability through feature-importance analysis, the presented AFE–DBE co-design framework can be beneficial to other machine-learning algorithms as well. The key principle is to align the AFE noise–power tradeoff with the noise and gain-error robustness of the downstream decoding algorithm, which can be generalized beyond the SPORF classifier employed here, as demonstrated in a recent study [40].

III. RECORDING AFE

In high-channel-count neural interfaces, direct-digitization architectures offer excellent area efficiency in scaled CMOS [25], [41]. Their digitally intensive nature allows for aggressive area scaling of digital logic, while supply-voltage scaling

reduces both digital and analog power [15]. In particular, incremental ADCs are especially attractive, offering superior area and power efficiency compared to their free-running counterparts [25], [42], due to much simpler decimation; with 1-bit internal quantization, the decimator reduces to a ripple counter [24].

Fig. 4 presents the proposed AFE channel, an ac-coupled transconductance operational amplifier-capacitor (OTA-C) incremental ADC with extended counting and a novel embedded IDAC in the feedback path. Each conversion starts with an incremental reset ($RST = 1$) to clear any memory effect. The ADC then runs 511 oversampling clock cycles to generate nine most significant bits (MSBs), followed by a one-cycle extended counting, yielding an overall 10-bit resolution. To ensure accurate residue evaluation for extended counting, the feedback is disabled ($EN_{FB} = 0$) in the first clock cycle of the incremental conversion, while the input is disabled ($EN_{IN} = 0$) during the extended counting phase. Input-reset switches, normally off, are turned on to short the OTA inputs to the common-mode voltage (V_B) only when the ADC output approaches saturation [25]. Implemented in 28-nm CMOS, the proposed AFE channel provides high input impedance, small area, low power, and robust rail-to-rail EDO rejection.

A. Ac Coupling With Pseudo-Resistor Biasing

In compact ac-coupled neural AFEs, small coupling capacitors (typically a few pF) are paired with pseudo-resistors to support the recording of very-low-frequency neural signals, often extending below 1 Hz. However, achieving such low-frequency recordings with small coupling capacitors requires extremely large resistors, which makes the pseudo-resistor a critical design element. Furthermore, the pseudo-resistor must maintain a stable input common-mode voltage for the OTA while introducing negligible noise and offset [43].

As shown in Fig. 5 (top), a pseudo-resistor can be efficiently implemented by operating a transistor in the OFF state, whose drain current is given by

$$I_{ds} = I_0 e^{\frac{V_{gs}-V_{th}}{nV_T}} \left(1 - e^{-\frac{V_{ds}}{V_T}}\right) \quad (1)$$

where V_{gs} is the gate-source voltage, V_{ds} the drain-source voltage, V_{th} the threshold voltage, V_T the thermal voltage, and I_0 and n are process-dependent parameters. Under desired operations with very small OTA input gate leakage (I_g) and μV -level input signals, $|V_{ds}| \ll V_T$, hence

$$1 - e^{-\frac{V_{ds}}{V_T}} \approx \frac{V_{ds}}{V_T}. \quad (2)$$

The corresponding small-signal resistance can be given by

$$r_{ds} = \left(\frac{\partial I_{ds}}{\partial V_{ds}}\right)^{-1} \propto e^{-\frac{V_{gs}}{nV_T}} = e^{\frac{V_B}{nV_T}} \quad (3)$$

indicating that the OTA gate bias voltage V_B modulates r_{ds} exponentially. Together with C_{in} , this resistance sets the high-pass corner [$f_c = 1/(2\pi r_{ds} C_{in})$] of the recording channel.

On the other hand, to ensure low noise, the dc resistance (R_{DS}) must be sufficiently large. Because its thermal noise




Fig. 5. AC-coupling with pseudo-resistor biasing (top) and simulated results (bottom), including small-signal resistance, noise power spectrum density (PSD), signal transfer function $|V_g/V_{in}|$, and input-referred noise (1–500 Hz).

$4kTR_{DS}$ is low-pass filtered with a corner frequency f_c , which equals the high-pass cutoff of the signal path. The root-mean-square (rms) noise integrated within the f_L - f_H signal band can be expressed as

$$V_{n,rms} = \sqrt{\frac{2kT}{\pi C} \cdot \left[\arctan\left(\frac{f_H}{f_c}\right) - \arctan\left(\frac{f_L}{f_c}\right) \right]} \quad (4)$$

which reduces to the well-known $(kT/C)^{1/2}$ noise when $f_L = 0$ and $f_H = \infty$. For $f_H \gg f_L \gg f_c$ (using $x \gg 1$, $\arctan(x) \approx \pi/2 - 1/x$)

$$V_{n,rms} \approx \sqrt{\frac{kT}{C}} \cdot \sqrt{\frac{2f_c}{\pi f_L}}. \quad (5)$$

Here, the attenuation factor $(2f_c/\pi f_L)^{1/2}$ accounts for the filtering effect and the specific bandwidth starting at f_L . Thus, increasing R_{DS} (i.e., reducing f_c) effectively reduces the integrated noise within the signal band.

Notably, the OFF-state dc resistance closely matches the small-signal ac resistance. This is because when $V_{ds} \ll V_T$, I_{ds} varies approximately linearly with V_{ds} , leading to

$$R_{DS} = \frac{V_{ds}}{I_{ds}} \approx \left(\frac{\partial I_{ds}}{\partial V_{ds}}\right)^{-1} = r_{ds}. \quad (6)$$

The simulation results, shown in Fig. 5 (bottom), indicate that while a V_B of 0 V is sufficient to set the f_c well below 1 Hz, a




Fig. 6. Multichannel direct-digitization recording architectures employing front-end OTA with (a) conventional IDAC and (b) proposed embedded IDAC.

substantially higher V_B of 150 mV, corresponding to an R_{DS} of $> 100 \text{ T}\Omega$, is required to suppress the $V_{n,rms}$ below $1.5 \mu\text{V}_{rms}$.

However, such a high resistance results in a large time constant $r_{ds}C_{in} > 100 \text{ s}$, leading to a very slow recovery from artifact-induced OTA saturation (e.g., motion or stimulation artifacts). To address this, a control signal is introduced to dynamically regulate the transistor's gate terminal (see SW_{in} in Fig. 4). During periods of saturation, the control signal turns the transistor on to enable rapid recovery. During normal conditions, it remains low to keep the transistor off and preserve the desired high OFF resistance [25], [44].

For a target linear range of 4 mV_{pp} at the AFE input, the biasing pseudo-resistor experiences a maximum voltage swing of 2 mV under unipolar recording configuration, where the reference input R_{ref} (see Fig. 4) is connected to a fixed bias voltage (ground in this design). Simulation results show that the pseudo-resistor induced total harmonic distortion (THD) remains below -130 dB . Under large artifacts (such as due to motion or stimulation) that drive the AFE input beyond its linear range, the THD can degrade quickly (e.g., -64 dB with 40 mV_{pp} artifacts). Nevertheless, the input reset logic pulls back the OTA input nodes (V_{gp} and V_{gn}) to the common mode voltage V_B whenever the AFE output approaches saturation. As a result, the voltage across the pseudo-resistor stays bounded within the AFE's linear input range (IR) and hence its voltage-dependent nonlinearity remains negligible.

Finally, the OTA input gate leakage (I_g) is another critical design consideration. When I_g flows through the off transistor, it creates a dc voltage drop that can lower the effective OFF resistance and may drive the OTA input outside its common-mode range, degrading performance or even compromising functionality. To address this, thick-oxide devices are typically employed due to their lower gate leakage. Unfortunately, based on the authors' experience, thick-oxide gate leakage is either poorly modeled or entirely omitted in most CMOS technologies, as modeling primarily focuses on thin-oxide devices for digital circuits. Nevertheless, a reasonable estimate can be obtained as follows.

Gate leakage in MOS transistors strongly depends on the gate-oxide thickness t_{ox} . For t_{ox} of a few nm, where direct tunneling dominates, the leakage current under a given device dimension and bias condition can be approximated by an

exponential relation [45]

$$I_g(t_{ox}) \approx K e^{(-\alpha t_{ox})} \quad (7)$$

where K and α are process-dependent parameters determined by the oxide stack and material properties. The data reported in [45] confirm this exponential dependence. Since t_{ox} values are specified in the process design kit (PDK) and thin-oxide gate leakage is typically well characterized, tunneling-based models can be employed to predict thick-oxide gate leakage. In addition, in the 28-nm technology used here, the high-k/metal-gate (HK/MG) stack lowers the leakage by over three orders of magnitude compared to silicon dioxide [45]. This leads to a total reduction of more than five orders of magnitude when t_{ox} is increased by $\sim 2 \text{ nm}$ (from thin to thick oxide). For the utilized device dimensions and bias conditions, a thin-oxide transistor exhibits $\sim 10 \text{ pA}$ of gate leakage, resulting in an estimated I_g of well below 0.1 fA for the thick-oxide devices. When incorporated into the simulations of Fig. 5 (top), this estimate shows negligible impact on overall voltage drop, transfer function, and noise performance.

B. Input Stage With Embedded IDAC

Incremental ADC-based neural readouts have demonstrated superior area and power efficiency [24], [25], [46], [47]. However, their multichannel implementations with conventional IDAC feedback present two critical challenges: 1) interchannel noise coupling caused by IDAC switching through shared bias lines; and 2) channel-to-channel gain variations [24].

Typically, the feedback IDAC is realized as either a current source or a current sink, rather than using both, to avoid mismatch between the two [24], [25]. To save power and area, and to simplify system-level configuration, channels often share OTA/IDAC bias lines. However, fast IDAC switching can inject noise onto these shared lines, degrading overall noise performance.

Hence, it is generally recommended to separate the OTA and IDAC bias lines. Such an N-channel example is shown in Fig. 6(a). The shared OTA bias line V_{GT} sets the OTA bias current I_{BIAS} , while a separate bias line V_{GB} establishes the feedback current I_{FB} . V_{gp} and V_{gn} denote the OTA input




Fig. 7. Simulated transconductance efficiency (g_m/I_D) and IR versus OTA bias current (I_{BIAS}).

signals, while D_p and D_n control the polarity of the feedback current. For simplicity, the remaining circuitry of the recording channel is not shown. Although the OTA and IDAC use separate biases, inter-IDAC coupling can still occur through the shared IDAC bias line V_{GB} . In addition, gate leakage currents I_{gt} and I_{gb} flowing through parasitic resistances R_{GT} and R_{GB} can introduce voltage drops, creating voltage gradients along the V_{GT} and V_{GB} lines. Notice that the two edge channels (CH_1 and CH_N) correspond to the extremes of gate leakage, with CH_1 showing the maximum (I_g) and CH_N the minimum (I_g). Furthermore, when the OTA tail transistors $M_{T1,...,N}$ and IDAC bias transistors $M_{B1,...,N}$ are widely distributed across the array, device-to-device mismatch increases and can deviate from the predictions of commonly used ABAB interdigitated layout models. These factors can result in significant channel-to-channel gain variations and increased noise levels. Large gain errors can degrade DBE performance for seizure detection, necessitating on-chip gain calibration [48].

To mitigate the noise coupling and achieve good channel uniformity without calibration, we propose an embedded IDAC seamlessly integrated into the OTA input pair, as illustrated in Fig. 6(b). This novel approach divides each input transistor into two smaller devices, with widths a and b (fixed length). The drain terminals of the transistors (b) are connected together, forming a current

$$I_{FB} = \frac{b}{a+b} I_{BIAS} \quad (8)$$

which is then steered to the positive or negative output based on the feedback control signals D_p and D_n . In this case, the input transconductance can be expressed as

$$g_m = \frac{g_m}{I_D} I_D = \frac{g_m}{I_D} \frac{a}{a+b} \frac{I_{BIAS}}{2} \quad (9)$$

and therefore the peak-to-peak IR is given by

$$IR = \frac{2I_{FB}}{g_m} = \frac{4b}{a} \frac{1}{g_m/I_D}. \quad (10)$$

Thus, the IR , and therefore the channel gain, is mainly set by the transistor width ratio (b/a) and, to first order, is independent of I_{BIAS} because the transconductance efficiency g_m/I_D remains nearly constant in weak inversion regions, which are typically chosen to maximize power efficiency. Since the width ratio can be accurately controlled with proper layout techniques, our embedded IDAC enables high channel-to-channel gain uniformity without calibration. Notably, this comes at the cost of reduced flexibility in adapting the IR




Fig. 8. Schematic of the front-end OTA with the embedded IDAC.

compared to a conventional IDAC, where the IR can be easily adjusted by trimming the IDAC reference current.

Because an IR of only a few mV is sufficient to accommodate the neural signals of interest, typical g_m/I_D values in weak inversions (~ 30 S/A) lead to $b \ll a$. Consequently, the embedded IDAC has a negligible impact on the operations of the input and tail transistors.

Fig. 7 presents the simulated g_m/I_D and the corresponding IR versus I_{BIAS} for $b/a = 1/38$. As expected, g_m/I_D remains largely constant (~ 32 S/A) over a wide I_{BIAS} range from 10 to 500 nA, thus ensuring stable IR (or channel gain). The g_m/I_D tends to drop slightly at higher bias current levels due to the minor increase in the inversion level. Interestingly, at very low bias currents, a slight decrease in g_m/I_D is also observed due to gate-induced drain leakage (GIDL) [49].

Although the embedded IDAC maintains good channel gain uniformity with minimal noise coupling, voltage drops along the OTA bias lines V_{GT} can still introduce spatial I_{BIAS} gradients and thus noise gradients, across channels. To address this, we implement local generations of I_{BIAS} , ensuring good matching before routing them to each input pair. Given the low current levels (up to 400 nA), even routing resistance (R_D) of a few k Ω introduces voltage drops of a few mV, which are negligible for circuit operation and performance.

For linearity, a similar analysis can be performed following the approach in [19], yielding the ADC output:

$$D_{out} \propto V_{in} - \frac{V_{in}^3}{12(nV_T)^2} \quad (11)$$

where V_{in} denotes the differential input signal. The behavior closely resembles that of a conventional IDAC paired with an input stage, where linearity is limited by the input differential pair operating in the weak inversion region. The dominant third-order harmonic distortion

$$HD_3 \approx \frac{1}{48} \left(\frac{V_m}{nV_T} \right)^2 \quad (12)$$

with V_m denoting the amplitude of the input signal [49].

Furthermore, since the switched transistors are much smaller than the always-on input transistors ($b \ll a$), the noise or

interference coupled onto the bias line during IDAC switching is minimal. In addition, our embedded IDAC eliminates the dummy current path typically employed in conventional IDACs to improve transient response during switching [24].

C. Front-End OTA

Fig. 8 shows the transistor-level schematic of the front-end OTA with the embedded IDAC. The width ratio a/b is set to 38 to achieve an IR of >3.5 mV_{PP} at $g_m/I_D \approx 32$ S/A (see Fig. 7). Each input device is further divided into two equally sized transistors to support the cross-coupled configuration that cancels the differential input during extended counting [46].

To reduce power, the input stage operates at a lower supply of 0.7 V, while the output stage is biased at one-tenth of the input current. The nominal I_{BIAS} is 400 nA and tunable for noise–power tradeoff. C_{gd} neutralization minimizes the effective gate parasitic capacitance to ~ 100 fF. With a compact 1.6 pF ac-coupling capacitor, the resulting capacitive division attenuates the signal by $<10\%$. Note that while the embedded IDAC introduces extra parasitic capacitance at the OTA input nodes, its impact is negligible because the IDAC devices (M_{FB}) are sized substantially smaller than the input transistors (M_{IP}) for a small IR of a few mV. However, when the IR is extended (e.g., to tens of mV to accommodate large artifacts), the added parasitics can become significant, thereby reducing effectiveness of the C_{gd} neutralization. The neutralization capacitors are implemented with high-density metal–oxide–metal (MOM) structures (~ 5 fF/ μm^2 , M1–M6), offering $\sim 7\times$ higher density than MOS capacitors. Nevertheless, this approach degrades the effectiveness of neutralization under different bias conditions, as will be shown later in the measurements. Furthermore, internal chopping ($f_{chp} = f_{chn} = 256$ kHz) suppresses offset and flicker noise, allowing the use of small load transistors (M_{LN} and M_{LP}). Finally, the load PMOS transistors are cross-coupled and self-biased by the OTA outputs, eliminating the requirement for an auxiliary common-mode feedback (CMFB) amplifier [46].

The simulated OTA input-referred noise is 3.94/5.28 μV_{rms} for an I_{BIAS} of 400/200 nA. Accounting for the noise contributions from the pseudo-resistors [$(2)^{1/2} \times 1.5$ μV_{rms}] and the capacitive attenuation ($\sim 10\%$) at the OTA input nodes, the AFE’s intrinsic input-referred noise is 4.92/6.26 μV_{rms} . As a comparison, the theoretical quantization noise of the AFE channel is 1.13 μV_{rms} , which is designed to be negligible compared to the intrinsic noise floors under various OTA bias conditions. Notably, the overall noise is designed with sufficient margin relative to the ~ 10 μV_{rms} tolerance to accommodate potential channel performance degradation and variations introduced in the measurements.

D. Comparator and Reconstruction

The single-bit internal quantizer is implemented with a strong-arm latch-based dynamic comparator [50], followed by an set–reset (SR) latch. This single-bit quantization enables a highly efficient decimation architecture, realized with a simple ripple counter. To minimize leakage power, all standard cell-based digital logic employs ultrahigh threshold voltage




Fig. 9. Channel schematic of the HV-compliant neurostimulator (anodic path omitted for simplicity).

(UHV) devices, which are sufficient to support the maximum operating frequency of 1.024 MHz.

IV. HV-COMPLIANT NEUROSTIMULATION

Electrical stimulation for epileptic seizure suppression typically requires current amplitudes of several hundred μA . Depending on the electrode–tissue interface impedance, this current can easily lead to voltage drops of several volts across the interface. HV technologies can support HV compliance to accommodate these large drops, but suffer from poor area and power efficiency when integrating the AFE and DBE [51]. LV CMOS with transistor stacking improves efficiency, yet remains constrained by the breakdown voltage of the parasitic junction diodes [14]. In the 28-nm CMOS technology employed in this work, the breakdown voltage is <10 V, which limits the maximum voltage compliance.

To decouple the tradeoff between hardware efficiency and HV compliance, the presented heterogeneous chipset architecture partitions the HV stimulator from the LV AFE and DBE. As shown in Fig. 9, each stimulation channel consists of a segmented 8-bit IDAC, a push–pull (PP) stage, and an output driver (the anodic path is omitted for simplicity). To improve power efficiency, only the output driver is implemented in HV BCD, while the IDAC and PP stage are realized in LV CMOS. The segmented IDAC, consisting of 4-bit binary and 4-bit unary segments, provides fine-grained current control (LSB ≈ 0.4 μA) and is implemented using cascode current sources. The subsequent PP stage converts the unipolar input current (I_{in}) into an anodic or cathodic current (I_A or I_C) depending on the operating phase. Finally, the HV driver delivers the programmed stimulation current to the electrodes.

To reduce power, the HV current driver operates with a gain of 8. A regulated cascode topology is employed to boost the output impedance. Bulk-driven feedback is proposed to improve the output swing and loop stability compared to the conventional gate-driven feedback architecture [52]. Because achieving a large output swing requires V_{DS2} to be small (<200 mV in this work), the conventional gate-driven approach, where V_{DS2} is applied to the gate terminal of M_{BN} ,




Fig. 10. Die microphotographs of the fabricated chipset, and summary of the chip dimensions, process nodes, and metal-stack specifications.




Fig. 11. Layout of the AFE channel and shared bias circuitry (per 32 channels), with corresponding area and power breakdown.

would require a large M_{BN} device to properly bias the branch indicated by the dashed gray box. The large device dimension increases parasitic capacitance and consequently degrades the phase margin of the impedance-boosting loop. In contrast, the proposed bulk-driven approach allows for a low V_{DS2} ($= V_{B2}$) while keeping M_{BN} compact, since the bias condition is primarily established by the gate bias V_D with an additional minor contribution from the bulk modulation through V_{B2} . The series resistor R_{B2} (17 k Ω) limits potentially large current flow through the bulk terminal during start-up. Under normal operating conditions, the voltage drop across R_{B2} is negligible because the bulk leakage I_{BL2} is small (<1 pA) at the targeted low values of V_{DS2} . In addition, the input branch employs a matched regulated cascode with replica bias, ensuring high mirror accuracy through matched V_{DS} ($V_{DS1} \approx V_{DS2}$). A small bleeder current ($I_{BLE} = 500$ nA) prevents the internal nodes from floating in the OFF state, which enables rapid transitions




Fig. 12. Measured channel output codes with dc component removed (top), and corresponding IRs (bottom) versus I_{BIAS} for (a) 400, (b) 300, and (c) 200 nA.

to the desired operating points when the driver is turned on. Operating from dual ± 5.5 V supplies, the driver enables zero dc bias across the stimulation electrodes.

V. MEASUREMENT RESULTS

A prototype chipset was fabricated using a hybrid process flow that combines 28-nm high-performance compact-plus (HPC+) and 180-nm HV-BCD technologies, with chip micrographs shown in Fig. 10. The 64-channel AFE is arranged in two groups, each sharing a bias circuitry. A zoomed-in layout of a single AFE channel and the shared bias circuitry (per 32 channels) is shown in Fig. 11. Due to well-proximity effect (WPE), the four edge-tail transistors $M_{T15,16,31,32}$ exhibit $\sim 2\%$ lower current than the others. The resulting AFE channel occupies an area of only $868 \mu\text{m}^2$, dominated by the OTA and input capacitors. Operating from 0.7 V (OTA input stage) and 0.9 V supplies, each channel consumes 204 nW, with the OTA input stage consuming $\sim 70\%$ of the power. The input-stage supply cannot be reduced below 0.7 V due to the constraint set by the input common-mode voltage ($V_B = 0.15$ V).

A. Electrical Characterizations

The 64-channel raw data, two 32-channel groups, were serialized using two serializers and streamed off-chip for analysis. In testing, only the first 12 channels from group 1 (channels 1–12) and the first 11 channels from group 2 (channels 33–43) consistently produced valid waveforms. The remaining channels exhibited varying levels of data corruption, such as unexpected glitches or jumps, likely caused by timing issues in the data serialization or in the subsequent off-chip acquisition chain. As a result, only channels 1–12 and 33–43 were considered in the subsequent analysis.




Fig. 13. Measured signal transfer function of the AFE channel.




Fig. 14. Measured input-referred noise voltage for a grounded input with internal chopping disabled (top left) and enabled (top right), and corresponding noise PSD (bottom).

The effectiveness of the proposed embedded IDAC was evaluated under three OTA bias currents: 400, 300, and 200 nA. As shown in Fig. 12 (top), for a 30-Hz 2-mV_{pp} input, the measured channel output codes (with dc components removed for better comparison) exhibit nearly identical amplitudes across all bias settings. The corresponding IR , presented in Fig. 12 (bottom), have mean values of 3.94, 3.99, and 4.05 mV, with standard deviations of 3.6%, 3.5%, and 3.6% ($n = 23$, 1σ), respectively. The IR variation appears random, with no observable gradient within each group. Notably, the two edge channels (1 and 33) exhibit the largest deviations, likely due to degraded matching of transistor width ratio b/a . Introducing a dummy channel at each edge could potentially mitigate this effect. A slight increase in IR is observed at lower bias currents, which contrasts with the simulated trend in Fig. 7(b). The deviation is attributed to variations in the parasitic gate capacitance of the input pair at reduced I_{BIAS} , which diminishes the effectiveness of C_{gd} neutralization. The resulting increase in capacitive attenuation marginally enlarges the IR . Nonetheless, the IR variations remain small, all at approximately 3.6%. Therefore, except for the comparison in this analysis, all measurement results are obtained with a bias current of 200 nA to reduce power consumption while maintaining acceptable noise performance.




Fig. 15. Measured channel output spectrum with 80- and 350-Hz input sinewaves, each applied at amplitudes of 0.2 and 2 mV_{pp}.




Fig. 16. Measured input reset events when the channel output code crosses the lower (63) or upper (960) threshold.

The measured signal transfer function of the AFE channel is shown in Fig. 13. The gain remains flat at low frequencies, with a 3-dB bandwidth of approximately 500 Hz. At higher frequencies, the response follows a sinc-shaped filtering profile, with notches at the Nyquist rate (1 kHz) and its harmonics, which reflects the inherent anti-aliasing of the continuous-time incremental conversion.

The measured input-referred noise voltage of all 23 channels is shown in Fig. 14 (top) with grounded input. Without internal chopping, several channels saturate and appear as flat traces near -2 and 2 mV, caused by large offsets introduced in the small PMOS/NMOS load transistors. Enabling internal chopping eliminates this saturation. The corresponding spectrum in Fig. 14 (bottom) shows that the internal chopping reduces the integrated noise (1–500 Hz) from 25.1 to $7.6 \mu\text{V}_{\text{RMS}}$. Since the OTA input pair is not chopped, flicker noise from the input transistors is not suppressed. Note that the measured noise level is $\sim 20\%$ higher than the simulated value ($6.26 \mu\text{V}_{\text{RMS}}$), which could be attributed to additional noise introduced by the measurement setup.

Fig. 15 shows the measured channel output spectrum for different sinewave inputs at 80 and 350 Hz. At an amplitude of 0.2 mV_{pp}, no harmonic distortion is observed. Increasing the amplitude to 2 mV_{pp} produces a clear second harmonic (HD₂) of the 80-Hz tone at 160 Hz, as well as an aliased HD₂ of the 350-Hz tone at 300 Hz. The strong HD₂ limits the THD to -51.4 dB. This second-order harmonic distortion originates from nonlinearities introduced by the single-ended input, where the reference input (see R_{ref} in Fig. 4) is tied to ground. Under such an input amplitude (2 mV_{pp}), the HD₃ predicted by (12) remains well below -90 dB.

Fig. 16 illustrates the occurrence of input reset events, triggered when a channel output crosses the predefined lower (63) or upper (960) threshold. This reset mechanism enables rapid recovery from saturation caused by any undesired artifacts.




Fig. 17. Measured driver output I - V characteristics at its maximum output current ($768 \mu\text{A}$).




Fig. 18. Measured driver output voltage with biphasic stimulation.

Notice that due to the channel's input-referred offset voltage and the switching-induced kT/C "noise", the output code immediately following each reset exhibits a deviation from its nominal mid-scale value of 512.

Fig. 17 shows the maximum output currents measured from the HV driver as a function of the output voltage, demonstrating stable current delivery over a wide voltage range from -5 to $+5$ V under ± 5.5 V supply rails. With biphasic stimulation, the measured output waveform ($6 \text{ k}\Omega + 470 \text{ nF}$ load) is presented in Fig. 18, achieving HV compliance up to 10 V. A residual voltage is clearly observed after stimulation due to the mismatch between the cathodic and anodic paths. While an H-bridge was not implemented in this design to minimize the chip pad count, future implementations can adopt this structure to reuse the same conduction path for both polarities, thereby minimizing the mismatch.

B. In vitro Validation

To further validate the functionality of the proposed neural recording chip in a biologically relevant environment, in vitro experiments were conducted with the setup shown in Fig. 19 (top). A 32-channel flexible μECoG array with $40 \times 40 \mu\text{m}$ electrodes (PEDOT:PSS) was immersed in a recording chamber filled with phosphate-buffered saline (PBS) to mimic physiological conditions. Pre-recorded ECoG signals were applied through an injection electrode (Ag/AgCl) using an arbitrary waveform generator (AWG). The array was positioned close to the injection electrode, while a reference electrode (Pt) was placed on the opposite side to provide a return path. The resulting saline potentials were sensed by the μECoG array and subsequently recorded by the ASIC. The recorded output waveforms are shown in Fig. 19 (bottom), alongside the original input signal for comparison, demonstrating close agreement. The recorded amplitudes are slightly reduced due to the voltage gradient in the saline between the injection and reference electrodes. Two reset events are




Fig. 19. Experimental setup for saline measurements (top), applied input signal (middle), and corresponding recorded output waveforms (bottom).

TABLE I
COMPARISON WITH OTHER NEUROMODULATION SoCs

Metric	Uehlin [53] JSSC20	Shin [14] JSSC22	Zhang [13] JSSC22	This Work
Architecture	Monolithic	Monolithic	Monolithic	Chipset
Technology [nm]	65	65	40	28/180 ^(a)
Channels (rec/stim)	64/4	256/4	16/1	64/4
HV compliance [V]	11	8	1	10 ^(b)
Stacking	YES	YES	NO	NO

^(a)HV BCD; ^(b)Up to 70 V with transistor stacking.

also triggered when the channel input crosses the -1.75 mV threshold.

C. Comparison to State-of-the-Art

Table I compares our proposed chipset with other bidirectional neuromodulation SoCs. In contrast to monolithic single-chip solutions, our presented heterogeneous chipset architecture achieves high area and power efficiency in the recording AFE, while simultaneously supporting HV compliance for the stimulation.

Table II summarizes the performance of the proposed neural AFE and compares it with state-of-the-art ECoG readouts (≤ 1 kHz bandwidth). Combining optimal architecture choices and the AFE-DBE co-design, our proposed AFE achieves the smallest per-channel area and the lowest per-channel power.

Furthermore, a novel embedded IDAC-based incremental $\Delta\Sigma$ ADC is proposed to address the limitations in conventional IDACs and open-loop recording architectures,

TABLE II
PERFORMANCE SUMMARY AND COMPARISON TO STATE-OF-THE-ART NEURAL AFEs

Metric	Kim [16] JSSC'18	Uehlin [18] TBioCAS'20	Shin [14] JSSC'22	Huang [19] JSSC'22	Wu [54] TBioCAS'24	Jiang [37] JSSC'25	This Work
Technology [nm]	65 (Bulk CMOS)	65 (Bulk CMOS)	65 (Bulk CMOS)	22 (FDSOI)	65 (Bulk CMOS)	180 (Bulk CMOS)	28 (Bulk CMOS)
Channel count	16	64	256	256	16	32	64
Channel topology	$\Delta\Sigma$	Multiplexed IA + SAR	Multiplexed IA + SAR	Multiplexed I- $\Delta\Sigma$	I- $\Delta\Sigma$	$\Delta\Sigma$	I- $\Delta\Sigma$
Input structure	AC-coupled w/ chopping	AC-coupled w/ CDS	AC-coupled w/ chopping	AC-coupled w/ chopping	AC-coupled w/ chopping	AC-coupled w/ chopping	AC-coupled^(c) w/o chopping
ADC Feedback DAC	CDAC	CDAC	CDAC	BDAC	CDAC	CDAC	Embedded IDAC
EDO Compensation DAC	CDAC	CDAC	CDAC	CDAC	N/A	CDAC	N/A
Supply [V]	0.8	0.5/2.5	1.2	0.6/0.8	1.0	1.8	0.7/0.9
Multiplexing	NO	64:1	256:4	256:16	NO	NO	NO
Area/Ch [mm ²]	0.024 ^(a)	0.0023	0.004	0.001	0.017 ^(a)	0.028 ^(a)	0.0009
Power/Ch [μ W]	0.8 ^(a)	2.98	1.51	1.61	5.56 ^(a)	5.39 ^(a)	0.204
Bandwidth [Hz]	1-500	1-1k	1-500	1-500	1-1k	1-500	1-500
IRN [μ V _{RMS}]	0.94	1.66	3.2	1.55	1	3.7	7.6
THD [%]	$\sim 0.05^{(b)}$ @31.7mV _{pp}	$\sim 1^{(b)}$ @0.6mV _{pp}	$\sim 0.1^{(b)}$ @2.5mV _{pp}	0.22 @2mV _{pp}	N/R	0.18 @10mV _{pp}	0.27 @2mV _{pp}
EDO Tolerance [mV _{pp}]	260	110	100	156.6	N/R	222	Rail-to-rail
Crosstalk [dB]	N/R	-92	-79.8/-72.8	-59 ^(d)	N/R	-68	-58 ^(d)
CMRR [dB]	81	76	70	98	N/R	70	61

^(a)Decimation filter off-chip and not included; ^(b)Estimated from measurements; ^(c)Chopping applied internally inside OTA; ^(d)Limited by PCB routing; N/R: Not reported; N/A: Not allowed.




Fig. 20. Comparison of channel power versus area for this work and prior neural AFEs with ECoG bandwidth (≤ 1 kHz).

namely the large channel-to-channel gain variations and the noise degradation caused by multichannel IDAC switching. The proposed embedded IDAC achieves high channel-gain uniformity without calibration and is robust against a wide range of bias current variations.

Fig. 20 presents a power-area scatter plot, comparing this work with prior neural AFEs with ECoG bandwidth (≤ 1 kHz). In general, time-division multiplexing (TDM) architectures achieve a smaller channel area. Compared to non-multiplexed designs (w/o TDM), this work achieves up to $5\times$ smaller area, while consuming over $3\times$ lower power, thanks to the employed algorithm-aware AFE optimization. Furthermore, our area is even smaller than those highly multiplexed designs, despite being implemented in an older CMOS technology compared to the leading 22-nm FDSOI design in [19]. Above all, prior work exclusively relies on input chopping or CDS to achieve very low noise, leading to significantly reduced input impedance and limited EDO rejection.

Finally, when our proposed design is scaled to 1000 channels, the projected resource requirements are only ~ 1 mm² in area and ~ 200 μ W in power, thus holding promise for the development of next-generation epilepsy monitoring and intervention systems with wide cortical coverage and high spatial resolution.

VI. CONCLUSION

In this article, we present a bidirectional neuromodulation chipset tailored for epileptic seizure management, with the key contributions summarized as follows.

- 1) A heterogeneous chipset architecture that delivers a highly area- and power-efficient neural AFE alongside a HV-compliant neurostimulator, overcoming the tradeoffs inherent to monolithic single-chip solutions.
- 2) An AFE and digital backend (DBE) co-design framework, which leverages algorithm-aware hardware optimization to significantly reduce the AFE's power and area, while maintaining decoding performance in the DBE.
- 3) A novel embedded current-steering IDAC that addresses key limitations of conventional IDACs, including noise degradation and channel-to-channel gain variations, thereby enhancing scalability and signal integrity in multichannel neural recording systems.
- 4) State-of-the-art performance, achieving the lowest per-channel power and the smallest per-channel area to date for neural AFEs with ECoG bandwidths (≤ 1 kHz).

Together, these contributions pave the way for more efficient and scalable closed-loop neuromodulation interfaces, offering a promising path toward compact, low-power implantable systems for the long-term management of epileptic seizures.

ACKNOWLEDGMENT

The authors would like to thank Zu-Yao Chang for his support with wire bonding, Brian Nanhekhan, Mustafeez Shah, and Bakr H. Abdelgalil for their help during chip testing, Ronald Wijermars for his help with HV-driver layout, and Peng Guo and Yi-Han Ou-Yang for the valuable technical discussions.

REFERENCES

- [1] V. R. Rao and J. D. Rolston, "Unearthing the mechanisms of responsive neurostimulation for epilepsy," *Commun. Med.*, vol. 3, no. 1, p. 166, Nov. 2023.
- [2] K. W. Scangos et al., "Closed-loop neuromodulation in an individual with treatment-resistant depression," *Nature Med.*, vol. 27, no. 10, pp. 1696–1700, Oct. 2021.
- [3] S. Stanslaski et al., "Sensing data and methodology from the adaptive DBS algorithm for personalized therapy in Parkinson's disease (ADAPT-PD) clinical trial," *NPJ Parkinson's Disease*, vol. 10, no. 1, Sep. 2024, Art. no. 174.
- [4] J. J. Ferrero et al., "Closed-loop electrical stimulation prevents focal epilepsy progression and long-term memory impairment," *Nat. Neurosci.*, vol. 28, no. 8, pp. 1753–1762, Aug. 2025.
- [5] M. T. Salam, J. L. Perez Velazquez, and R. Genov, "Seizure suppression efficacy of closed-loop versus open-loop deep brain stimulation in a rodent model of epilepsy," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 24, no. 6, pp. 710–719, Jun. 2016.
- [6] K. Abdelhalim, H. M. Jafari, L. Kokarovtseva, J. L. P. Velazquez, and R. Genov, "64-channel UWB wireless neural vector analyzer SoC with a closed-loop phase synchrony-triggered neurostimulator," *IEEE J. Solid-State Circuits*, vol. 48, no. 10, pp. 2494–2510, Oct. 2013.
- [7] M. A. Bin Altaf, C. Zhang, and J. Yoo, "A 16-channel patient-specific seizure onset and termination detection SoC with impedance-adaptive transcranial electrical stimulator," *IEEE J. Solid-State Circuits*, vol. 50, no. 11, pp. 2728–2740, Nov. 2015.
- [8] H. Kassiri et al., "Rail-to-rail-input dual-radio 64-channel closed-loop neurostimulator," *IEEE J. Solid-State Circuits*, vol. 52, no. 11, pp. 2793–2810, Nov. 2017.
- [9] C.-H. Cheng et al., "A fully integrated 16-channel closed-loop neural-prosthetic CMOS SoC with wireless power and bidirectional data telemetry for real-time efficient human epileptic seizure control," *IEEE J. Solid-State Circuits*, vol. 53, no. 11, pp. 3314–3326, Nov. 2018.
- [10] G. O'Leary et al., "A recursive-memory brain-state classifier with 32-channel track-and-zoom $\delta^2\Sigma$ s ADCs and charge-balanced programmable waveform neurostimulators," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 296–298.
- [11] G. O'Leary et al., "26.2 A neuromorphic multiplier-less bit-serial weight-memory-optimized 1024-tree brain-state classifier and neuromodulation SoC with an 8-channel noise-shaping SAR ADC array," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2020, pp. 402–404.
- [12] Y. Wang et al., "A closed-loop neuromodulation chipset with 2-level classification achieving 1.5-vpp CM interference tolerance, 35-dB stimulation artifact rejection in 0.5ms and 97.8%-sensitivity seizure detection," *IEEE Trans. Biomed. Circuits Syst.*, vol. 15, no. 4, pp. 802–819, Aug. 2021.
- [13] M. Zhang, L. Zhang, C.-W. Tsai, and J. Yoo, "A patient-specific closed-loop epilepsy management SoC with one-shot learning and online tuning," *IEEE J. Solid-State Circuits*, vol. 57, no. 4, pp. 1049–1060, Apr. 2022.
- [14] U. Shin et al., "NeuralTree: A 256-channel 0.227- μ J/class versatile neural activity classification and closed-loop neuromodulation SoC," *IEEE J. Solid-State Circuits*, vol. 57, no. 11, pp. 3243–3257, Nov. 2022.
- [15] C. M. Lopez and X. Huang, "Circuits and architectures for neural recording interfaces," in *Biomedical Electronics, Noise Shaping ADCs, and Frequency References: Advances in Analog Circuit Design 2022*, P. Harpe, A. Baschiroto, and K. A. A. Makinwa, Eds., Cham, Switzerland: Springer, 2023, pp. 45–57.
- [16] C. Kim, S. Joshi, H. Courellis, J. Wang, C. Miller, and G. Cauwenberghs, "Sub- μ V_{rms}-noise sub- μ W/channel ADC-direct neural recording with 200-mV/ms transient recovery through predictive digital autoranging," *IEEE J. Solid-State Circuits*, vol. 53, no. 11, pp. 3101–3110, Nov. 2018.
- [17] M. Sharma, H. J. Strathman, and R. M. Walker, "Verification of a rapidly multiplexed circuit for scalable action potential recording," *IEEE Trans. Biomed. Circuits Syst.*, vol. 13, no. 6, pp. 1655–1663, Dec. 2019.
- [18] J. P. Uehlin, W. A. Smith, V. R. Pamula, S. I. Perlmutter, J. C. Rudell, and V. S. Sathe, "A 0.0023 mm²/ch. Delta-encoded, time-division multiplexed mixed-signal ECoG recording architecture with stimulus artifact suppression," *IEEE Trans. Biomed. Circuits Syst.*, vol. 14, no. 2, pp. 319–331, Apr. 2020.
- [19] X. Huang et al., "Actively multiplexed μ ECoG brain implant system with incremental- $\Delta\Sigma$ ADCs employing bulk-DACs," *IEEE J. Solid-State Circuits*, vol. 57, no. 11, pp. 3312–3323, Nov. 2022.
- [20] X. Huang, X. Yang, A. Lodi, C. Van Hoof, G. Gielen, and C. M. Lopez, "A 3072-channel neural readout IC with multiplexed two-step incremental-SAR conversion and bulk-DAC-based EDO compensation in 22 nm FDSOI," in *Proc. IEEE Symp. VLSI Technol. Circuits (VLSI Technol. Circuits)*, Jun. 2024, pp. 1–2.
- [21] R. Müller, S. Gambini, and J. M. Rabaey, "A 0.013 mm², 5 μ W, DC-coupled neural signal acquisition IC with 0.5 V supply," *IEEE J. Solid-State Circuits*, vol. 47, pp. 232–243, Jan. 2012.
- [22] W. A. Smith, J. P. Uehlin, S. I. Perlmutter, J. C. Rudell, and V. S. Sathe, "A scalable, highly-multiplexed delta-encoded digital feedback ECoG recording amplifier with common and differential-mode artifact suppression," in *Proc. Symp. VLSI Circuits*, Jun. 2017, pp. C172–C173.
- [23] M. Jang et al., "A 1024-channel 268-nW/pixel 36 $\times$ 36 μ m²/channel data-compressive neural recording IC for high-bandwidth brain-computer interfaces," *IEEE J. Solid-State Circuits*, vol. 59, no. 4, pp. 1123–1136, Apr. 2024.
- [24] D. De Dorigo et al., "Fully immersible subcortical neural probes with modular architecture and a delta-sigma ADC integrated under each electrode for parallel readout of 144 recording sites," *IEEE J. Solid-State Circuits*, vol. 53, no. 11, pp. 3111–3125, Nov. 2018.
- [25] X. Huang et al., "A compact, low-power analog front-end with event-driven input biasing for high-density neural recording in 22-nm FDSOI," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 69, no. 3, pp. 804–808, Mar. 2022.
- [26] R. Müller et al., "A minimally invasive 64-channel wireless μ ECoG implant," *IEEE J. Solid-State Circuits*, vol. 50, no. 1, pp. 344–359, Jan. 2015.
- [27] N. Even-Chen et al., "Power-saving design opportunities for wireless intracortical brain-computer interfaces," *Nature Biomed. Eng.*, vol. 4, no. 10, pp. 984–996, Aug. 2020.
- [28] S. R. Nason et al., "A low-power band of neuronal spiking activity dominated by local single units improves the performance of brain-machine interfaces," *Nature Biomed. Eng.*, vol. 4, no. 10, pp. 973–983, Jul. 2020.
- [29] X. Huang, H. Yassin, B. L. Alcazar, A. Akhondi, and D. G. Muratore, "A closed-loop neuromodulation chipset with 0.0009 mm²-0.36 μ W/Ch recording frontend and 0.075 mm²-6.76 μ W seizure classification backend," in *Proc. IEEE Symp. VLSI Technol. Circuits*, Jun. 2025, pp. 1–2.
- [30] E. F. Chang, "Towards large-scale, human-based, mesoscopic neurotechnologies," *Neuron*, vol. 86, no. 1, pp. 68–78, Apr. 2015.
- [31] M. Shaeri, U. Shin, A. Yadav, R. Caramellino, G. Rainer, and M. Shoaran, "A 2.46-mm² miniaturized brain-machine interface (MiBMI) enabling 31-class brain-to-text decoding," *IEEE J. Solid-State Circuits*, vol. 59, no. 11, pp. 3566–3579, Nov. 2024.
- [32] N. Pu et al., "A 36-nW electrocardiogram anomaly detector based on a 1.5-bit non-feedback delta quantizer for always-on cardiac monitoring," *IEEE Trans. Biomed. Circuits Syst.*, vol. 18, no. 3, pp. 648–661, Jun. 2024.
- [33] B. C. Johnson et al., "An implantable 700 μ W 64-channel neuromodulation IC for simultaneous recording and stimulation with rapid artifact recovery," in *Proc. Symp. VLSI Circuits*, Japan, Jun. 2017, pp. C48–C49.
- [34] J. Huang et al., "A 0.01-mm² mostly digital capacitor-less AFE for distributed autonomous neural sensor nodes," *IEEE Solid-State Circuits Lett.*, vol. 1, no. 7, pp. 162–165, Jul. 2018.
- [35] W. Jiang, V. Hokhikyan, H. Chandrakumar, V. Karkare, and D. Markovic, "A $\pm$ 50-mV linear-input-range VCO-based neural-recording front-end with digital nonlinearity correction," *IEEE J. Solid-State Circuits*, vol. 52, no. 1, pp. 173–184, Jan. 2017.
- [36] J.-S. Bang, H. Jeon, M. Je, and G.-H. Cho, "A 6.5 μ W 92.3db-dr biopotential-recording front-end with 360 mv_{pp} linear input range," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2018, pp. 239–240.
- [37] X. Jiang et al., "A bi-directional neural interface chip with 32-channel 83-dB DR CTDSM-based recording using FIRDAC with pre-emptive ELD compensation," *IEEE J. Solid-State Circuits*, early access, Jun. 4, 2025, doi: [10.1109/JSSC.2025.3571637](https://doi.org/10.1109/JSSC.2025.3571637).
- [38] A. Burrello, L. Cavigelli, K. Schindler, L. Benini, and A. Rahimi, "Laelaps: An energy-efficient seizure detection algorithm from long-term human iEEG recordings without false alarms," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Burrel, Italy, Mar. 2019, pp. 752–757.

- [39] T. M. Tomita et al., "Sparse projection oblique randomer forests," *J. Mach. Learn. Res.*, vol. 21, no. 104, pp. 1–39, 2015.
- [40] A. Pandey, D. Vaish, I.-T. Lin, and R. Müller, "Neural recording power optimization through machine learning guided resolution reconfiguration," 2025, *arXiv:2510.22924*.
- [41] X. Yang et al., "An AC-coupled 1st-order Δ - Δ readout IC for area-efficient neural signal acquisition," *IEEE J. Solid-State Circuits*, vol. 58, no. 4, pp. 949–960, Apr. 2023.
- [42] S. Wang et al., "A compact chopper stabilized Δ - Δ neural readout IC with input impedance boosting," *IEEE Open J. Solid-State Circuits Soc.*, vol. 1, pp. 67–78, 2021.
- [43] X. Huang and D. G. Muratore, "Recording front-end electronics for large-scale implantable brain–computer interfaces: A design perspective," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, Apr. 2025, pp. 1–8.
- [44] H. Gao et al., "HermesE: A 96-channel full data rate direct neural interface in 0.13 μm CMOS," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 1043–1055, Apr. 2012.
- [45] E. N. Shauly, "CMOS leakage and power reduction in transistors and circuits: Process and layout considerations," *J. Low Power Electron. Appl.*, vol. 2, no. 1, pp. 1–29, Jan. 2012.
- [46] D. Wendler et al., "A 0.0046-mm² two-step incremental delta-sigma analog-to-digital converter neuronal recording front end with 120-mV_{pp} offset compensation," *IEEE J. Solid-State Circuits*, vol. 58, no. 2, pp. 439–450, Feb. 2023.
- [47] R. Willaredt, C. Grandauer, D. De Dorigo, M. Kuhl, and Y. Manoli, "A 0.00175 mm² and 9.62 μW channel direct-digitization front end with EDO compensation for neural probes with a spatial resolution of 35 μm ," *IEEE J. Solid-State Circuits*, vol. 60, no. 7, pp. 2330–2340, Jul. 2025.
- [48] H. Kassiri et al., "Battery-less tri-band-radio neuro-monitor and responsive neurostimulator for diagnostics and treatment of neurological disorders," *IEEE J. Solid-State Circuits*, vol. 51, no. 5, pp. 1274–1289, May 2016.
- [49] P. G. Jespers and B. Murmann, *Systematic Design of Analog CMOS Circuits*. Cambridge, U.K.: Cambridge Univ. Press, 2017.
- [50] B. Razavi, "The design of a comparator [the analog mind]," *IEEE Solid State Circuits Mag.*, vol. 12, no. 4, pp. 8–14, Apr. 2020.
- [51] B. C. Johnson et al., "An implantable 700- μW 64-channel neuromodulation IC for simultaneous recording and stimulation with rapid artifact recovery," in *Proc. Symp. VLSI Circuits*, Jun. 2017, pp. C48–C49.
- [52] J. Charthad et al., "A mm-sized wireless implantable device for electrical stimulation of peripheral nerves," *IEEE Trans. Biomed. Circuits Syst.*, vol. 12, no. 2, pp. 257–270, Apr. 2018.
- [53] J. P. Uehlin et al., "A single-chip bidirectional neural interface with high-voltage stimulation and adaptive artifact cancellation in standard CMOS," *IEEE J. Solid-State Circuits*, vol. 55, no. 7, pp. 1749–1761, Jul. 2020.
- [54] J. Wu et al., "A low-noise low-power 0.001Hz–1 kHz neural recording system-on-chip with sample-level duty-cycling," *IEEE Trans. Biomed. Circuits Syst.*, vol. 18, no. 2, pp. 263–273, Apr. 2024.



Xiaohua Huang (Member, IEEE) received the B.Sc. degree in electrical engineering from Jilin University, Changchun, China, in 2015, the M.Sc. degree in electrical engineering from Tsinghua University, Beijing, China, in 2018, and the Ph.D. degree in electrical engineering from KU Leuven (in collaboration with IMEC), Leuven, Belgium, in 2023.

Since 2023, he has been a Post-Doctoral Researcher with Delft University of Technology, Delft, The Netherlands. From 2014 to 2018, he conducted research on the integrated circuit design

of phase-locked loop (PLL)-based clock generation and modulation systems for wireline applications. His Ph.D. Research focuses on analog and mixed-signal integrated circuit design and system development for high-density neural recording interfaces, exploring the heterogeneous integration of CMOS and thin-film transistor (TFT) technologies. His Post-Doctoral Research has expanded this scope to include on-chip neural signal processing, and focused ultrasound neuromodulation.



Hoda Yassin (Student Member, IEEE) received the B.Sc. and M.Sc. degrees in electronics and communication engineering from Aswan University, Aswan, Egypt, in 2021 and 2024, respectively. She is currently pursuing the Ph.D. degree with the Microelectronics Department, Delft University of Technology, Delft, The Netherlands.


From 2021 to 2022, she worked as a Research Associate at Aswan Power Electronic Applications Research Centre (APEARC), Aswan. In 2022, she was an Interned at Si-Vision Academy, Cairo, Egypt.

From 2022 to 2023, she joined the Information Technology Institute (ITI), Aswan. Her research focuses on analog and mixed-signal circuits and systems design for high-density closed-loop neuromodulation systems.



Beatriz Lafuente Alcazar received the B.Sc. degree in biomedical engineering from the Universitat Politècnica de València, Valencia, Spain, in 2022, and the M.Sc. degree (cum laude) in biomedical engineering from Delft University of Technology, Delft, The Netherlands, in 2024, where she is currently pursuing the Ph.D. degree with the Microelectronics Department.

Her current research interests include digital circuit design for machine-learning-based neural interfaces.



Arash Akhondi (Graduate Student Member, IEEE) received the B.Sc. degree in electrical engineering from the University of Tehran, Tehran, Iran, in 2015, and the M.Sc. degree in electrical engineering from Sharif University of Technology, Tehran, in 2018. He is currently pursuing the Ph.D. degree in bioelectronics section with the Department of Microelectronics, Delft University of Technology, Delft, The Netherlands.

From 2018 to 2022, he worked as a Digital Signal Processing and a FPGA Engineer, where he gained experience in signal processing implementation and digital IC design. His research focuses on massively parallel neural interfaces, including on-chip data compression, low-power spike sorting, and real-time neural signal processing.



Dante G. Muratore (Senior Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from the Politecnico di Torino, Turin, Italy, in 2012 and 2013, respectively, and the Ph.D. degree in microelectronics from the Integrated Microsystems Laboratory, University of Pavia, Pavia, Italy, in 2017.

From 2015 to 2016, he was a Visiting Scholar with Microsystems Technology Laboratories, Massachusetts Institute of Technology, Cambridge, MA, USA. From 2016 to 2020, he was a Post-Doctoral

Fellow with Stanford University, Stanford, CA, USA. Since 2020, he has been an Assistant Professor with the Department of Microelectronics, Delft University of Technology, Delft, The Netherlands, leading the Smart Brain Interfaces Group. His group investigates hardware and system solutions for high-bandwidth brain–machine interfaces that can interact with the nervous system at natural resolution. They contribute solutions for massively parallel bidirectional interfaces, on-chip neural signal processing, and wireless power and data transfer.

Dr. Muratore was a recipient of the Wu Tsai Neurosciences Institute Interdisciplinary Scholar Award and the ERC Consolidator Grant.